arXiv:1704.04362v1 [cs.DS] 14 Apr 2017

FAST MONTE CARLO ALGORITHMS FOR TENSOR OPERATIONS*

DAVOUD ATAEE TARZANAGH† AND GEORGE MICHAILIDIS‡

Abstract. We consider scalable *randomized* algorithms for many common tensor operations, including tensor multiplication, low-rank decomposition and nuclear norm minimization. Such operations arise in a number of modern day applications of tensors in machine learning and signal processing. Specifically, we introduce polynomial time algorithms that employ a small number of lateral and/or horizontal slices of the underlying 3-rd order tensor, that offer relative error guarantees for the quality of the solutions. All results can easily be extended to higher order tensors. The proposed methods are illustrated on selected real data sets.

Key words. slices selection, tensor, circulant algebra, low rank decomposition, nuclear norm minimization

AMS subject classifications. 15A69, 15A23, 68W20, 65F30

1. Introduction.

Tensors have been an object of intense study by different communities due to their diverse applications, including chemometrics [31], psychometrics [21], and image and signal processing [16, 17, 25, 30, 38].

A key operation is that of tensor decomposition due to its suitability in a number of data analysis tasks, including compression and classification [7, 32], latent variable modeling [1], etc. Another one is that of tensor multiplication [20]. Both operations become computationally expensive for very large size tensors. Further, motivated by recent developments in robust factorization of noisy data matrices and completion problems, there have been extensions in the literature to the tensor case [23, 30]. Hence, the broad objective of this paper is to develop fast Monte Carlo algorithms for tensor multiplication and decomposition and building on these developments use them for noisy tensor completion and robust factorization. Another important goal of this work is to provide the necessary probabilistic recovery guarantees. A key technical development that we leverage to establish the results is the use of *circulant algebra* embedding.

Next, we summarize some previous work on tensor decompositions. Suppose that $\boldsymbol{\mathcal{A}}$ is a $n_1 \times n_2 \times n_3$ (3-way) data tensor. Popular decompositions in the literature for obtaining a compressed representation of the data is the Core Decomposition (CP) [6], the TUCKER decomposition [15] and the t-SVD one [3, 18, 19]; the latter provides the "best" tubal rank-r approximation to $\boldsymbol{\mathcal{A}}$, as measured with respect to any unitary invariant tensor norm.

One major shortcoming of these three classical decompositions is their brittleness with respect to severely corrupted or outlying data points. To that end, a number of approaches have been developed to recover a low-rank tensor representation from data subject to noise and corrupted entries. We focus on two instances of the problem based on the t-SVD algorithm: (i) noisy tensor completion, i.e., recovering a low-rank tensor from a small subset of noisy entries, and (ii) noisy robust tensor factorization, i.e., recovering a low-rank tensor from corrupted data through noise and/or outliers of

*Submitted.

Funding: This work was supported in part by NSF grants DMS-1545277, DMS-1632730, NIH grant1R01-GM1140201A1 and by the UF Informatics Institute.

†Department of Mathematics & UF Informatics Institute, University of Florida, Gainesville, FL, tarzanagh@ufl.edu.

‡Department of Statistics & UF Informatics Institute, University of Florida, Gainesville, FL, gmichail@ufl.edu.

This manuscript is for review purposes only.

arbitrary magnitude [25, 38]. These two classes of tensor factorization problems have attracted significant interest in the research community [2, 4, 30, 38]. In particular, convex formulations of noisy tensor factorization have been shown to exhibit strong theoretical recovery guarantees and a number of algorithms has been developed for solving them [38, 30, 25].

It is frequently mentioned that (noisy) tensor factorization, despite its numerous advantages, also exhibits a number of drawbacks listed below:

- These methods [3, 19, 18, 38, 30, 25] are inherently sequential and all rely on the repeated and costly computation of t-SVD factors, Discrete Fourier Transform (DFT), and its inverse (IDFT) that limit the scalability of the algorithms.
- The basis tensor vectors resulting from t-SVD have little concrete meaning, which makes it very difficult for users to interpret the obtained results. For instance, the vector [(1/2) age - $(1/\sqrt{2})$ height + (1/2)income], being one of the significant uncorrelated factors from a data set of people's features is not easily interpretable (see discussion in [11]). Kuruvilla et al in [22] have also claimed: "it would be interesting to try to find basis vectors for all experiment vectors, using actual experiment vectors and not artificial bases that offer little insight."
- The t-SVD decomposition for sparse tensors does not preserve sparsity in general, which for large size tensors leads to excessive computations and storage requirements. Hence, it is important to compute low-rank tensor factorizations that preserve such structural properties of the original data tensor.

To address these challenges, we study scalable *randomized* tensor multiplication (rt-product algorithm) and tensor factorization (rt-project) operations and extend the matrix CX and CUR type decompositions [9, 10, 11] to third order tensors using a circulant algebra embedding. Therefore, we develop a basic algorithm (t-CX), as well as a more general algorithm (t-CUR) with relative error guarantees. Moreover, we provide a CUR tensor nuclear norm minimization (CUR t-NN) method which solves noisy tensor factorization using a small number of lateral and/or horizontal slices of the underlying tensor. To the best of our knowledge, this is the first work that provides such decompositions with relative-error recovery guarantees. Further, the proposed algorithms for these learning tasks have polynomial time complexity.

The remainder of the paper is organized as follows. In section 2, we review some relevant mathematical concepts including the tensor circulant algebra, basic definitions and theorems of tensors, and the t-SVD decomposition based on the t-product concept. In section 3, we present a randomized tensor-product and provide the key results on the quality of the approximation. In section 4, we provide our main results, i.e., the randomized tensors decomposition with relative error guarantees and introduce the concept of slice selection and projection. In section 5, we introduce, evaluate and analyze our proposed algorithm (CUR t-NN) for large scale noisy tensor decomposition. Results are provided in section 6. Finally, we conclude in section 7.

The detailed proofs of the main results obtained are provided in an Appendix.

2. Mathematical notation and tensor basics. Next, we review relevant mathematical concepts including tensor SVD, basic definitions and operations and other technical results of tensors, that are used throughout the paper.

Tensor indexing. We denote tensors by boldface Euler script letters, e.g., $\mathcal{X}$, matrices by boldface capital letters, e.g., X, vectors by lowercase letters, e.g., x. The order of a tensor is the number of dimensions (also refereed to as ways

or modes). In this work, we deal with 3-way tensors.

Fibers and slices [35]. A *fiber* of tensor $\mathcal{X}$ is a one-dimensional array defined by fixing two indices. $\mathcal{X}_{:jk}$ is the (j,k)-th *column* fiber, $\mathcal{X}_{i:k}$ is the (i,k)-th *row* fiber, and $\mathcal{X}_{ij:}$ is the (i,j)-th *tube* fiber. A slice of tensor $\mathcal{X}$ is a two-dimensional array defined by fixing one index. $\mathcal{X}_{i::}$ is the i-th *horizontal* slice, $\mathcal{X}_{:j:}$ is the j-th *lateral* slice, and $\mathcal{X}_{::k}$ is the k-th *frontal* slice. For convenience, $\mathcal{X}_{::k}$ is written as $\mathcal{X}_k$. The vectorization of $\mathcal{X}$ is denoted by $vec(\mathcal{X})$. For a 3-way tensor $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, we denote its (i,j,k)-th entry as $\mathcal{X}_{ijk}$.

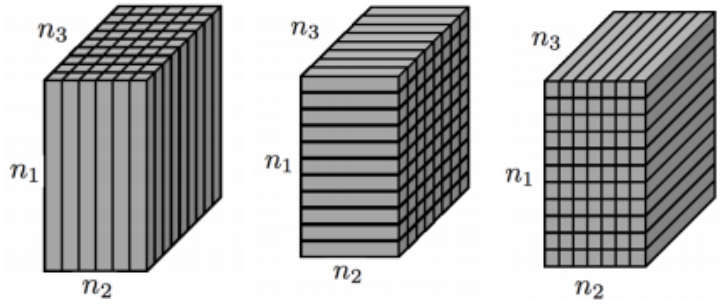

(a) column fibers; $\mathcal{X}_{:jk}$ (b) row fibers; $\mathcal{X}_{i:k}$ (c) tube fiber; $\mathcal{X}_{ij:}$

Fig. 1: Fibers of an $n_1 \times n_2 \times n_3$ tensor $\mathcal{X}$.

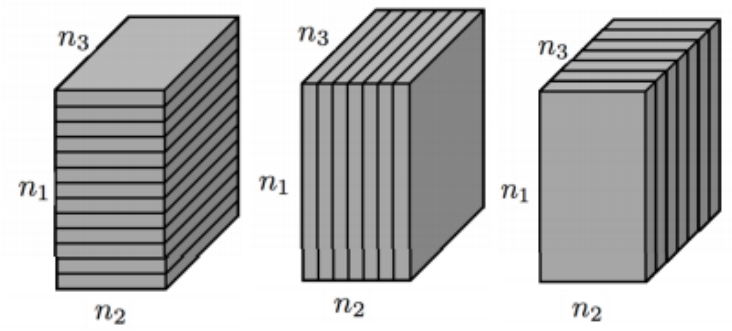

(a) horizontal slices: $\mathcal{X}_{i::}$ (b) lateral slices: $\mathcal{X}_{:j:}$ (c) frontal slices: $\mathcal{X}_{::k}$

Fig. 2: Slices of an $n_1 \times n_2 \times n_3$ tensor $\mathcal{X}$.

Norms. We denote the ℓ_1 norm as $\|\mathcal{X}\|_1 := \sum_{ijk} |x_{ijk}|$, the infinity norm as $\|\mathcal{X}\|_\infty := \max_{ijk} |x_{ijk}|$, the Frobenius norm as $\|\mathcal{X}\|_F := \sqrt{\sum_{ijk} |x_{ijk}|^2}$, and the spectral norm as $\|\mathcal{X}\| := \sigma_{\max}(\mathcal{X})$, where $\sigma_{\max}$ denotes the largest singular value of $\mathcal{X}$. The above norms reduce to the vector or matrix norms if $\mathcal{X}$ is a vector or a matrix.

Operators. For $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, by using the Matlab command `fft`, we denote $\hat{\mathcal{X}}$ as the result of applying the Discrete Fourier Transform (DFT) on $\mathcal{X}$ along

the 3-rd dimension, i.e., $\hat{\mathcal{X}} = \texttt{fft}(\mathcal{X}, [], 3)$. In the same fashion, one can also compute $\mathcal{X}$ from $\hat{\mathcal{X}}$ via the Inverse Discrete Fourier Transform (IDFT), using $\texttt{ifft}(\hat{\mathcal{X}}, [], 3)$. In particular, we denote $\hat{X}$ as a block diagonal matrix with each block on diagonal being the frontal slice $\hat{X}_{::k}$ of $\hat{\mathcal{X}}$.

DEFINITION 2.1. *[19]**(Tensor product)**. Given two tensors $\mathcal{Z} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and $\mathcal{X} \in \mathbb{R}^{n_2 \times n_4 \times n_5}$. The t-product $\mathcal{Z} * \mathcal{X}$ is the $n_1 \times n_4 \times n_3$ tensor ,*

$$\mathcal{C} = \mathcal{Z} * \mathcal{X} = \text{fold}\left(\text{circ}\left(\mathcal{Z}\right)\right).\text{unfold}\left(\mathcal{X}\right), \tag{1}$$

where

$$\text{circ}\left(\mathcal{Z}\right) := \begin{bmatrix} \mathcal{Z}_{::1} & \mathcal{Z}_{::n_3} & \dots & \mathcal{Z}_{::2} \\ \mathcal{Z}_{::2} & \mathcal{Z}_{::1} & \dots & \mathcal{Z}_{::3} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{Z}_{::n_3} & \mathcal{Z}_{::n_3-1} & \dots & \mathcal{Z}_{::1} \end{bmatrix},$$

and

$$\text{unfold}\left(\mathcal{X}\right) := \begin{bmatrix} X_{::1} \\ X_{::2} \\ \vdots \\ X_{::n_3} \end{bmatrix}, \qquad \text{fold}\left(\text{unfold}\left(\mathcal{X}\right)\right) = \mathcal{X}.$$

Because the circular convolution of two tube fibers can be computed by the DFT, the t-product can be alternatively computed in the Fourier domain, as shown in Algorithm 1.

Algorithm 1 t-product $\mathcal{C} = \mathcal{Z} * \mathcal{X}$ in the Fourier domain [19]

Input: $\mathcal{Z} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$; $\mathcal{X} \in \mathbb{R}^{n_2 \times n_4 \times n_3}$
$\hat{\mathcal{Z}} \leftarrow \texttt{fft}(\mathcal{Z}, [], 3)$;
$\hat{\mathcal{X}} \leftarrow \texttt{fft}(\mathcal{X}, [], 3)$;
for $k = 1, \dots, n_3$ **do**
  $\hat{\mathcal{C}}_{::k} = \hat{\mathcal{Z}}_{::k} \hat{\mathcal{X}}_{::k}$;
end for
$\mathcal{C} \leftarrow \texttt{ifft}(\hat{\mathcal{C}}, [], 3)$;
return $\mathcal{C}$

DEFINITION 2.2. *[19]**(Conjugate transpose)**. The conjugate transpose of a tensor $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is tensor $\mathcal{X}^* \in \mathbb{R}^{n_2 \times n_1 \times n_3}$ obtained by conjugate transposing each of the frontal slices and then reversing the order of transposed frontal slices 2 through n_3.*

DEFINITION 2.3. *[19]**(Identity tensor)**. The identity tensor $\mathcal{I} \in \mathbb{R}^{n \times n \times n_3}$ is the tensor whose first frontal slice is the $n \times n$ identity matrix, and whose other frontal slices are all zeros.*

DEFINITION 2.4. *[19]**(Orthogonal tensor).** A tensor $\mathcal{Q} \in \mathbb{R}^{n \times n \times n_3}$ is orthogonal if it satisfies $\mathcal{Q}^* * \mathcal{Q} = \mathcal{Q} * \mathcal{Q}^* = \mathcal{I}$.*

DEFINITION 2.5. *[19] **(F-diagonal Tensor)**. A tensor is called $f-$diagonal if each of its frontal slices is a diagonal matrix.*

THEOREM 2.6. *[19]**(t-SVD)**. Let* $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$. *Then, it can be factored as*

$$\mathcal{X} = \mathcal{U} * \Sigma * \mathcal{V}^T. \tag{2}$$

where $\mathcal{U} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, $\mathcal{V} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ *are orthogonal and* $\Sigma \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ *is a f-diagonal tensor.*

Note that t-SVD can be efficiently computed based on the matrix SVD in the Fourier domain. This is based on a key property that the block circulant matrix can be mapped to a block diagonal matrix in the Fourier domain, i.e.

$$(F_{n_3} \otimes I_{n_1}) \cdot \operatorname{circ}(\mathcal{X}) \cdot (F_{n_3}^{-1} \otimes I_{n_2}) = \hat{X}, \tag{3}$$

where F_{n_3} denotes the $n_3 \times n_3$ DFT matrix and $\otimes$ denotes the Kronecker product.

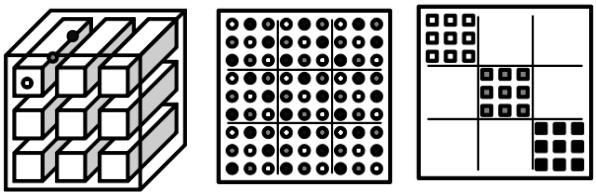

$\mathcal{X}$. circ($\mathcal{X}$). `fft`($\mathcal{X}$).

Fig. 3: The sequence of transformations in `fft` operation. Given a tensor $\mathcal{X}$, we convert it into a circulant matrix by circ ($\mathcal{X}$). In the third figure, we diagonalize each circulant using the `fft` operation.

DEFINITION 2.7. *[19]**(Tensor multi and tubal rank)**. The tensor multi-rank of* $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ *is a vector* $\upsilon \in \mathbb{R}^{n_3}$ *with its* $i-th$ *entry as the rank of the i-th frontal slice of* $\hat{\mathcal{X}}$, $r_i = \operatorname{rank}(\hat{\mathcal{X}}_{i::})$. *The tensor tubal rank, denoted as* $r = \operatorname{rank}_t(\mathcal{X})$, *is defined as the number of nonzero singular tubes of* Σ, *where* Σ *is from the t-SVD of* $\mathcal{X} = \mathcal{U} * \Sigma * \mathcal{V}^T$. *That is,*

$$r = card\{i : \Sigma_{ii:} \neq 0\} = \max_i r_i. \tag{4}$$

where card denotes the cardinality of a set.

The tensor tubal rank has some properties of the matrix rank, e.g.

$$\operatorname{rank}_t(\mathcal{X} * \mathcal{Z}) \leq \min(\operatorname{rank}_t(\mathcal{X}), \operatorname{rank}_t(\mathcal{Z})).$$

DEFINITION 2.8. *[30]**(Tensor nuclear norm)**. The tensor nuclear norm of a tensor* $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, *denoted as* $\|\mathcal{X}\|_{\circledast}$, *is defined as the average of the nuclear norm of all frontal slices of* $\hat{\mathcal{X}}$, *i.e.*

$$\|\mathcal{X}\|_{\circledast} := \frac{1}{n_3} \sum_{k=1}^{n_3} \|\hat{\mathcal{X}}_{::k}\|_*. \tag{5}$$

The above tensor nuclear norm is defined in the Fourier domain. It is closely related to the nuclear norm of the block circulant matrix in the original domain. Indeed,

$$\begin{aligned}\|\boldsymbol{\mathcal{X}}\|_{\circledast} &= \frac{1}{n_3}\sum_{k=1}^{n_3}\|\hat{\boldsymbol{\mathcal{X}}}_{::k}\|_* = \frac{1}{n_3}\|\hat{X}\|_* \\ &= \frac{1}{n_3}\|(F_{n_3}\otimes I_{n_1})\cdot \operatorname{circ}(\boldsymbol{\mathcal{X}})\cdot(F_{n_3}^{-1}\otimes I_{n_2})\|_* \\ &= \frac{1}{n_3}\|\operatorname{circ}(\boldsymbol{\mathcal{X}})\,\|_*\end{aligned}$$

The above relationship gives an equivalent definition of the tensor nuclear norm in the original domain. So the tensor nuclear norm is the nuclear norm (with a factor $1/n_3$) of a new matricization (block circulant matrix) of a tensor.

DEFINITION 2.9. *[19]**(**Inverse of tensor**)**. The inverse of* $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n\times n\times n_3}$ *is written as* $\boldsymbol{\mathcal{A}}^{-1}$ *satisfying*

$$\boldsymbol{\mathcal{A}}^{-1} * \boldsymbol{\mathcal{A}} = \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{A}}^{-1} = \boldsymbol{\mathcal{I}}, \tag{6}$$

where $\boldsymbol{\mathcal{I}}$ *is the identity tensor of size* $n \times n \times n_3$*.*

DEFINITION 2.10. *[37] **(Standard tensor basis).** The **lateral basis** $\mathring{\mathbf{e}}_i$, is of size $n_1 \times 1 \times n_3$ with only one entry equaling 1 and the rest equaling zero, in which the nonzero entry 1 will only appear at the first frontal slice of $\mathring{\mathbf{e}}_i$. Normally its transpose $\mathring{\mathbf{e}}_i^\top$ is called **horizontal basis**. The other standard tensor basis is called **tube basis** $\dot{\mathbf{e}}_i$, is a tensor of size $1 \times 1 \times n_3$ with one entry equaling to 1 and rest equaling to 0. Figure 4 illustrates these bases.*

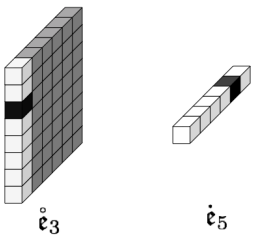

Fig. 4: The lateral basis $\mathring{\mathbf{e}}_3$ and tube basis $\dot{\mathbf{e}}_5$. The black cubes are 1, gray and white cubes are 0. The white cubes stand for the potential entries that could be 1 [37].

Since not all tensors can be recovered from data sets with missing entries and/or large outliers, recent theoretical advances have studied sufficient conditions for accurate recovery. Most prevalent among these are tensor *coherence conditions*, which limit the extent to which the singular vectors of a tensor are correlated with the standard tensor basis. Next, we define three standard notions of coherence for tensors.

DEFINITION 2.11. *[37]**(**Tensor* μ_0*-coherence**)**. Let* $\boldsymbol{\mathcal{U}} \in \mathbb{R}^{n_1\times r\times n_3}$*, and* $\boldsymbol{\mathcal{V}} \in$

$\mathbb{R}^{n_2\times r\times n_3}$, *satisfy* $\boldsymbol{\mathcal{U}}^T * \boldsymbol{\mathcal{V}} = \boldsymbol{\mathcal{I}}$. *Then the* μ_0*-coherence of* $\boldsymbol{\mathcal{U}}$ *and* $\boldsymbol{\mathcal{V}}$ *is:*

$$\mu_0(\boldsymbol{\mathcal{U}}) := \frac{n_1 n_3}{r} \max_{1\le i\le n_1} \|\boldsymbol{\mathcal{U}}^T * \mathring{\mathbf{e}}_i\|_F^2,$$
$$\mu_0(\boldsymbol{\mathcal{V}}) := \frac{n_2 n_3}{r} \max_{1\le j\le n_2} \|\boldsymbol{\mathcal{V}}^T * \mathring{\mathbf{e}}_j\|_F^2,$$

DEFINITION 2.12. *(**Tensor** μ_1**-coherence**). For a tensor* $\boldsymbol{\mathcal{L}} \in \mathbb{R}^{n_1\times n_2\times n_3}$ *assume that* $rank_t(\boldsymbol{\mathcal{L}}) = r$. *Then, the* μ_1*-coherence of* $\boldsymbol{\mathcal{L}}$ *is defined*

$$\mu_1(\boldsymbol{\mathcal{L}}) := \frac{n_1 n_2 n_3^2}{r} \|\boldsymbol{\mathcal{U}} * \boldsymbol{\mathcal{V}}^T\|_\infty^2,$$

where $\boldsymbol{\mathcal{L}}$ *has the skinny t-SVD* $\boldsymbol{\mathcal{L}} = \boldsymbol{\mathcal{U}} * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{V}}^T$.

DEFINITION 2.13. *(**Tensor** (μ, r)**-coherence**). For any* $\boldsymbol{\mu} > 0$, *we will call a tensor* $\boldsymbol{\mathcal{L}}$ (μ, r)*-coherent if*

$$rank_t(\boldsymbol{\mathcal{L}}) = r,$$
$$\max\{\mu_0(\boldsymbol{\mathcal{U}}), \mu_0(\boldsymbol{\mathcal{V}})\} \le \mu,$$
$$\mu_1(\boldsymbol{\mathcal{L}}) \le \mu.$$

Note that the standard tensor coherence is much weaker than the *matrix weak coherence condition* for each frontal slice of $\hat{\boldsymbol{\mathcal{X}}}$ [37]. Hence in our analysis of the proposed algorithms, we will use the standard tensor coherence condition.

3. Approximate tensor multiplication. In this section, we present a randomized tensor-product and provide the key results on the quality of the approximation. Given two tensors $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$, using the Definition 2.1, the t-product may be written in the following manner

$$\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} = \sum_{t=1}^{n_2} \boldsymbol{\mathcal{A}}_{:t:} * \boldsymbol{\mathcal{B}}_{t::}, \tag{7}$$

where $*$ denotes the tensor product.

It can be easily shown that the left hand side of (7) is equivalent to the block multiplication and then summation in the Fourier domain. When a tensor multiplication is formulated as (7), we can develop a randomized algorithm to approximate the product $\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}}$.

3.1. The new algorithm. Algorithm 2 takes as input two tensors $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$, positive integer $c \le n_2$, and a probability distribution $\{p_i\}_{i=1}^{n_2}$ over $[n_2]$. It returns as output two tensors $\boldsymbol{\mathcal{C}}$ and $\boldsymbol{\mathcal{R}}$, where the lateral slices of $\boldsymbol{\mathcal{C}}$ are a small number of sampled and rescaled slices of $\boldsymbol{\mathcal{A}}$ and similarly the horizontal slices of $\boldsymbol{\mathcal{R}}$ constitute a small number of sampled and rescaled slices of $\boldsymbol{\mathcal{B}}$. Formally speaking, assume that at most c lateral slices (in expectation) of $\boldsymbol{\mathcal{A}}$ are chosen by including the i-th lateral slice of $\boldsymbol{\mathcal{A}}$ in $\boldsymbol{\mathcal{C}}$ with probability $p_i = \min\{1, cp_i\}$. Then, define the sampling tensor $\boldsymbol{\mathcal{S}} \in \mathbb{R}^{n_2\times n_2\times n_3}$ to be the zero-one tensor where $S_{ii:} = 1$ if the i-th slice is chosen and $\boldsymbol{\mathcal{S}}_{ij} = 0$ otherwise, and define the rescaling tensor $\boldsymbol{\mathcal{D}} \in \mathbb{R}^{n_2\times c\times n_3}$ to be the tensor with $\boldsymbol{\mathcal{D}}_{ij:} = 1/\sqrt{cp_j}$ if $i-1$ of the previous slice have been chosen and $\boldsymbol{\mathcal{D}}_{ij:} = 0$ otherwise. Similarly, the `rt-product` sample and rescale the corresponding horizontal slices of the tensor $\boldsymbol{\mathcal{B}}$. In both of these cases, $\boldsymbol{\mathcal{C}} = \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}$ is an $n_1 \times c \times n_3$ tensor consisting of sampled and rescaled copies of the lateral slices of $\boldsymbol{\mathcal{A}}$, and $\boldsymbol{\mathcal{R}} = (\boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}})^\top * \boldsymbol{\mathcal{B}} = \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{B}}$ is a $c \times n_2 \times n_3$ tensor consisting of sampled and rescaled copies of the horizontal slices of $\boldsymbol{\mathcal{B}}$. For the case of $n_3 = 1$, the algorithm chooses column-row pairs in Algorithm 2 and is identical to the algorithm of [10].

Algorithm 2 (rt-product), a fast Monte-Carlo algorithm for approximate tensor multiplication.

Input: $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, $\boldsymbol{\mathcal{B}} \in \mathbb{R}^{n_2 \times n_4 \times n_3}$, $p_i \geq 0, i \in [n_2]$ s.t. $\sum_{i \in [n_2]} p_i = 1$, positive integer $c \leq n_2$.
Initialize $\boldsymbol{\mathcal{S}} \in \mathbb{R}^{n_2 \times n_2 \times n_3}$ and $\boldsymbol{\mathcal{D}} \in \mathbb{R}^{n_2 \times c \times n_3}$ to the all zeros tensors.
$t = 1$;
for $i = 1, \ldots, n_2$ **do**
  Pick i with probability $\min\{1, cp_i\}$;
  if i is picked **then**
    $\boldsymbol{\mathcal{S}}_{it1} = 1, \quad \boldsymbol{\mathcal{S}}_{it2:n_3} = 0$;
    $\boldsymbol{\mathcal{D}}_{tt1} = 1/\min\{1, \sqrt{cp_i}\}, \quad \boldsymbol{\mathcal{D}}_{tt2:n_3} = 0$;
    $t = t + 1$
  end if
end for
$\boldsymbol{\mathcal{C}} = \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}, \quad \boldsymbol{\mathcal{R}} = \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{B}}$;
$\hat{\boldsymbol{\mathcal{C}}} \leftarrow \texttt{fft}(\boldsymbol{\mathcal{C}}, [], 3), \quad \hat{\boldsymbol{\mathcal{R}}} \leftarrow \texttt{fft}(\boldsymbol{\mathcal{R}}, [], 3)$;
for $k = 1, \ldots, n_3$ **do**
  $\hat{\boldsymbol{\mathcal{Z}}}_{::k} = \hat{\boldsymbol{\mathcal{C}}}_{::k} \hat{\boldsymbol{\mathcal{R}}}_{::k}$;
end for
$\boldsymbol{\mathcal{Z}} \leftarrow \texttt{ifft}(\hat{\boldsymbol{\mathcal{Z}}}, [], 3)$;
return $\boldsymbol{\mathcal{Z}}$

3.2. Running time of rt-product. The `rt-product` is straightforward to implement, is computationally efficient, has small RAM requirements, and is well suited for large scale problems. Indeed, the `rt-product` algorithm can be implemented without storing the tensors $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and $\boldsymbol{\mathcal{B}} \in \mathbb{R}^{n_2 \times n_4 \times n_3}$ and uses $O(\max(n_1, n_4) c n_3 \log(n_3))$ flops to transform to the Fourier domain and $O(c(n_1 + n_2 + n_4)n_3)$ flops to construct $\boldsymbol{\mathcal{C}} \in \mathbb{R}^{n_1 \times c \times n_3}$ and $\boldsymbol{\mathcal{R}} \in \mathbb{R}^{c \times n_4 \times n_3}$, where

$$\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} \sim \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}. \tag{8}$$

It is worth mentioning that we do not need the storage of the sampling and resclaing tensors $\boldsymbol{\mathcal{S}}$ and $\boldsymbol{\mathcal{D}}$ for our numerical implementation.

3.3. Theoretical analysis . Next, we provide the main result on the quality of the approximation obtained from Algorithm 2 which states that, under appropriate assumptions (8) holds. The most interesting of these assumptions is that the sampling probabilities used to randomly sample the lateral slices of $\boldsymbol{\mathcal{A}}$ and the corresponding horizontal slices of $\boldsymbol{\mathcal{B}}$ are *non-uniform* and depend on the product of the norms of the lateral slices of $\boldsymbol{\mathcal{A}}$ and/or the corresponding horizontal slices of $\boldsymbol{\mathcal{B}}$.

We consider two examples of nonuniform sampling probabilities:

i. If we would like to use the information from both tensors $\boldsymbol{\mathcal{A}}$ and $\boldsymbol{\mathcal{B}}$, we consider sampling probabilities $\{p_i\}_{i=1}^{n_2}$ such that

$$p_i \geq \beta \frac{\|\boldsymbol{\mathcal{A}}_{:i:}\|_F \|B_{i::}\|_F}{\sum_{i=1}^{n_2} \|\boldsymbol{\mathcal{A}}_{:i:}\|_F \|\boldsymbol{\mathcal{B}}_{i::}\|_F}, \qquad \beta \in (0, 1]. \tag{9}$$

ii. If, for example, only information about $\boldsymbol{\mathcal{A}}$ is easily available. We use sampling probabilities $\{p_i\}_{i=1}^{n_2}$ such that

$$p_i \geq \beta \frac{\|\boldsymbol{\mathcal{A}}_{:i:}\|_F}{\|\boldsymbol{\mathcal{A}}\|_F}, \qquad \beta \in (0, 1]. \tag{10}$$

The following theorem gives the main result for Algorithm 2, and generalizes Theorem 6 in [11] to third order tensors.

THEOREM 3.1. *Suppose* $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, $\mathcal{B} \in \mathbb{R}^{n_2 \times n_4 \times n_5}$, *and* $c \leq n_2$. *Under the notion of Algorithm 2, If the sampling probabilities* $\{p_i\}_{i=1}^{n_2}$ *used by the algorithm are of the form (9) or (10), then*

$$\mathbf{E}[\|\mathcal{A} * \mathcal{B} - \mathcal{C} * \mathcal{R}\|_F] \leq \frac{1}{\sqrt{\beta c}} \|\mathcal{A}\|_F \|\mathcal{B}\|_F. \tag{11}$$

COROLLARY 3.2. *Given a tensor* $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ *with* $r \geq \text{rank}_t(\mathcal{A})$, *an error tolerance* $\epsilon \in (0, 1]$, *and a failure probability* $\delta \in (0, 1]$, *define sampling probabilities* p_i *satisfying* (10). *Under the notion of Algorithm 2, choose*

$$c \geq \xi r \log(4r/(\beta\delta))/(\beta\epsilon^2),$$

where ξ *is a positive constant and let* $\mathcal{C} \in \mathbb{R}^{n_1 \times c \times n_3}$ *and* $\mathcal{R} = \mathcal{C}^\top$ *be the corresponding subtensors of* $\mathcal{A}$. *Then,*

$$\|\mathcal{A} * \mathcal{A}^\top - \mathcal{C} * \mathcal{D} * \mathcal{D} * \mathcal{C}^\top\|_F^2 \leq \frac{\epsilon}{2} \|\mathcal{A}\|_F^2,$$

with probability at least $1 - \delta$.

4. Approximate Tensor Low Rank Decomposition. Motivated by numerous applications in which the data may be modeled by a variable subscripted by three or more indices, we develop a tensor-based extension of the matrix CX and CUR decomposition. Our main algorithms result in computing low-rank tensor approximations that are explicitly expressed in terms of a small number of slices of the input tensor.

4.1. Linear algebra with tensors: free submodules. The set of complex numbers $\mathbb{C}$ with standard scalar addition and multiplication is a field and $\mathbb{C}^{n_3}$ forms a vector space over this field. However, as pointed out in [17, 19], the set of tubes $\mathbb{C}^{1 \times 1 \times n_3}$ equipped with the tensor product form a ring with unity. A module over a ring can be thought of as a generalization of a vector space over a field, where the corresponding scalars are the elements of the ring. In linear algebra over a ring, the analog of a subspace is a *free submodule*. Our algorithm relies on submodules, and the following theorem.

THEOREM 4.1. *[3] The set of slices*

$$\Upsilon := \{\mathcal{X}_{:j:} \mid \quad \mathcal{X}_{:j:} \in \mathbb{C}^{n_1 \times 1 \times n_3} \quad \textit{and} \quad j \in [n_2]\}, \tag{12}$$

forms a free module over the set of tubes $\mathbb{C}^{1 \times 1 \times n_3}$.

Using Theorem 4.1, Υ has a *basis* so that any element of it can be written as a "t-linear combination" of elements of a set of basis slices. By "t-linear combination", we mean a sum of slices multiplied, with the t-product, by coefficients from $\mathbb{C}^{1 \times 1 \times n_3}$; we provide the following definition (see also, Figure 5).

DEFINITION 4.2. *(**Slice-wise linear independence**) The slices in a subset* $\Lambda = \{\mathcal{X}_{:1:}, \ldots, \mathcal{X}_{:n:}\}$ *of* Υ *are said to be* linearly dependent, *if there exist a finite number of distinct slices* $\mathcal{X}_{:1:}, \ldots, \mathcal{X}_{:m:}$ *in* Λ, *and tubes* $\mathcal{C}_{11:}, \ldots, \mathcal{C}_{mm:}$, *not all zero, such that*

$$\sum_{j=1}^{m} \mathcal{X}_{:j:} * \mathcal{C}_{jj:} = O, \tag{13}$$

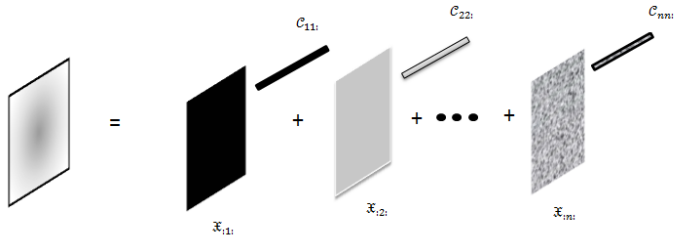

Fig. 5: An element of Υ generated by t-linear combinations of slices $\mathfrak{X}_{:1:}, \mathfrak{X}_{:2:}, \ldots, \mathfrak{X}_{:n:}$ with coefficients $\mathfrak{C}_{11:}, \mathfrak{C}_{22:}, \ldots, \mathfrak{C}_{nn:}$.

where O denotes the lateral slice of all zeros.

The slices in a subset $\Lambda = \{\mathfrak{X}_{:1:}, \ldots, \mathfrak{X}_{:n:}\}$ of Υ are said to be linearly independent *if the equation*

$$\sum_{j=1}^{n} \mathfrak{X}_{:j:} * \mathfrak{C}_{jj:} = O,$$

can only be satisfied by all zero tubes $\mathfrak{C}_{jj:}$, $j = 1, \ldots n$.

4.2. Slice-based low rank approximation: tensor CX decomposition. Using the concept of free submodules and Definition 4.2, we introduce for the first time the notion of *slice selection* and *projection*.

DEFINITION 4.3. *Let $\mathfrak{C} \in \mathbb{R}^{n_1 \times c \times n_3}$ and $\mathfrak{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$. Then, the tensor project (t-project) $\Pi_{\mathfrak{C}}(\mathfrak{X})$ is an $n_1 \times n_2 \times n_3$ tensor,*

$$\Pi_{\mathfrak{C}}(\mathfrak{X}) := \mathfrak{C} * \mathfrak{C}^{+} * \mathfrak{X}, \tag{14}$$

where $*$ denotes t-product 1, $\Pi_{\mathfrak{C}}(\mathfrak{X})$ is the projection of $\mathfrak{X}$ onto the subspace spanned by the lateral slices of $\mathfrak{C}$, and $\mathfrak{C}^{+}$ is the Moore-Penrose generalized inverse of the tensor $\mathfrak{C}$.

Because the circulant convolution of two tube fibers can be computed by the DFT the t-project can be alternatively computed in the Fourier domain, as shown in Algorithm 3.

Algorithm 3 t-project $\Pi_{\mathfrak{C}}(\mathfrak{X})$ computation in the Fourier domain

Input: $\mathfrak{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$; $\mathfrak{C} \in \mathbb{R}^{n_1 \times c \times n_3}$.
$\hat{\mathfrak{X}} \leftarrow \texttt{fft}(\mathfrak{X}, [], 3)$;
$\hat{\mathfrak{C}} \leftarrow \texttt{fft}(\mathfrak{C}, [], 3)$;
for $k = 1, \ldots, n_3$ **do**
  $\hat{\mathfrak{Z}}_{::k} = \hat{\mathfrak{C}}_{::k} \hat{\mathfrak{C}}^{+}_{::k} \hat{\mathfrak{X}}_{::k}$;
end for
$\mathfrak{Z} \leftarrow \texttt{ifft}(\hat{\mathfrak{Z}}, [], 3)$;
return $\Pi_{\mathfrak{C}}(\mathfrak{X}) = \mathfrak{Z}$

DEFINITION 4.4. *(**t-CX**). Let $\mathfrak{X}$ be an $n_1 \times n_2 \times n_3$ tensor. For any given tensor $\mathfrak{C}$, an $n_1 \times c \times n_3$ tensor whose lateral slices consist of c lateral slices of the tensor $\mathfrak{X}$,*

the $n_1 \times n_2 \times n_3$ *tensor*

$$\Pi_{\mathcal{C}}(\mathcal{X}) = \mathcal{C} * \mathcal{C}^{\dagger} * \mathcal{X},$$

is a slice-based tensor approximation to $\mathcal{X}$.

Several things should be noted about this definition.

- The choice of $c \ll n_2$ depends on the application. For example, c could be constant, independent of the dimension of tensor, logarithmic in the size of n_2, or a large constant factor less than n_2.
- The t-CX decomposition expresses each of the slices of $\mathcal{X}$ in terms of a linear combination of basis slices (see, Figure 5), each of which is an actual lateral slice $\mathcal{X}$. Thus, a t-CX decomposition provides a low-rank approximation to the original tensor, although one with structural properties that are different than those provided by the t-SVD.
- Given a set of lateral slices $\mathcal{C}$, the approximation $\Pi_{\mathcal{C}}(\mathcal{X}) = \mathcal{C} * \mathcal{C}^{\dagger} * \mathcal{X}$ is the "best" approximation to $\mathcal{X}$ in the sense that

$$\|\mathcal{X} - \mathcal{C} * (\mathcal{C}^{+} * \mathcal{X})\|_F = \min_{\mathcal{Y} \in \mathbb{R}^{n_1 \times c \times n_3}} \|\mathcal{X} - \mathcal{C} * \mathcal{Y}\|_F.$$

Next, we describe the algorithm and provide a theorem, from which our first main result, the relative error for t-CX, follows. Algorithm 4 takes as input an $n_1 \times n_2 \times n_3$ tensor $\mathcal{A}$, a tubal rank parameter r, and an error parameter ϵ. It returns as output an $n_1 \times c \times n_3$ Tensor $\mathcal{C}$ consisting of a small number of slices of $\mathcal{A}$. Let $\mathcal{L} = \mathcal{U} * \Sigma * \mathcal{V}^{\top}$, be a rank-$r$ approximation of tensor $\mathcal{A}$. Algorithm 4 uses the sampling probabilities

$$p_i \geq \frac{\beta}{r n_3} \|\hat{\mathcal{V}}_{:i:}^{\top}\|_F^2, \qquad \forall i \in [n_2], \qquad \beta \in (0,1], \tag{15}$$

to sample a small number of lateral slices of $\mathcal{A}$, where $\hat{\mathcal{V}} = \texttt{fft}(\mathcal{V})$. Note also that, we define rescaling and sampled tensors $\hat{\mathcal{D}}$ and $\hat{\mathcal{S}}$ in the Fourier domain. This leads to a significant reduction of time complexity of the fast Fourier transform and its inverse. For the case of $n_3 = 1$, the algorithm selects column-row pairs in Algorithm 4 and is identical to the algorithm of [11].

Using Corollary 3.2, we now establish a tensor version of Lemma 1 of [11].

LEMMA 4.5. *Let* $\epsilon \in (0,1]$ *and* $\mathcal{L} = \mathcal{U} * \Sigma * \mathcal{V}^{\top}$ *be a rank-r approximation of* $\mathcal{A}$. *Define* $\Gamma = (\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D})^{+} - (\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D})^{\top}$. *Under the notion of Algorithm 4, if* $c \geq \xi r \log(4r/(\beta\delta))/(\beta\epsilon^2)$, *where* ξ *is a positive constant and* $\delta \in (0,1]$. *Then, with probability at least* $1 - \delta$,

$$\begin{aligned} \text{rank}_t(\mathcal{V}^{\top} * \mathcal{S}) &= \text{rank}_t(\mathcal{V}) = \text{rank}_t(\mathcal{L}), \\ \|\Gamma\| &= \|\Sigma^{-1}_{\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D}} - \Sigma_{\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D}}\|, \\ (\mathcal{L} * \mathcal{S} * \mathcal{D})^{+} &= (\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D})^{+} \Sigma^{-1} \mathcal{U}^{\top}, \\ \|\Sigma^{-1}_{\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D}} - \Sigma_{\mathcal{V}^{\top} * \mathcal{S} * \mathcal{D}}\| &\leq \epsilon/\sqrt{2}. \end{aligned}$$

Using Lemma 4.5, we provide an improved sampling complexity for the randomized ℓ_2 regression of [11]. The proof follows along similar lines to the proof of Theorem 5 of [11].

PROPOSITION 4.6. *Let* $\epsilon \in (0,1]$ *and* $\mathcal{L} = \mathcal{U} * \Sigma * \mathcal{V}^{\top}$ *be a rank-r approximation of* $\mathcal{A}$. *If* $c \geq \xi r \log(4r/(\beta\delta))/(\beta\epsilon^2)$, *where* ξ *is a positive constant and* $\delta \in (0,1]$, *then with probability at least* $1 - \delta - 0.2$,

$$\|\mathcal{A} - \mathcal{A} * \mathcal{S} * \mathcal{D} * (\mathcal{L} * \mathcal{S} * \mathcal{D})^{+} * \mathcal{L}\|_F \leq (1+\epsilon)\|\mathcal{A} - \mathcal{A} * \mathcal{L}^{+} * \mathcal{L}\|_F.$$

Algorithm 4 (t-CX), a fast Monte-Carlo algorithm for tensor low rank approximation.

Input: $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, $p_i \geq 0$, a rank parameter r, positive integer $c \leq n_2$.
$\hat{\boldsymbol{\mathcal{A}}} \leftarrow \texttt{fft}(\boldsymbol{\mathcal{A}}, [], 3)$;
Compute sampling probabilities p_i for all $i \in [n_2]$ given by (15);
Initialize $\hat{\mathbf{S}}$ and $\hat{\boldsymbol{\mathcal{D}}}$ to the all zeros tensors.
$t = 1$;
for $i = 1, \ldots, n_2$ **do**
  Pick i with probability $\min\{1, cp_i\}$;
  if i is picked **then**
    $\hat{\mathbf{S}}_{it1} = 1, \quad \hat{\mathbf{S}}_{it2:n_3} = 0$;
    $\hat{\boldsymbol{\mathcal{D}}}_{tt1} = 1/\min\{1, \sqrt{cp_i}\}, \quad \hat{\boldsymbol{\mathcal{D}}}_{tt2:n_3} = 0$;
    $t = t + 1$
  end if
end for
for $k = 1, \ldots, n_3$ **do**
  $\hat{\boldsymbol{\mathcal{Z}}}_{::k} = \hat{\boldsymbol{\mathcal{A}}}_{::k}\hat{\mathbf{S}}_{::k}\hat{\boldsymbol{\mathcal{D}}}_{::k}(\hat{\boldsymbol{\mathcal{A}}}_{::k}\hat{\mathbf{S}}_{::k}\hat{\boldsymbol{\mathcal{D}}}_{::k})^{+}\hat{\boldsymbol{\mathcal{A}}}_{::k}$;
end for
$\boldsymbol{\mathcal{Z}} \leftarrow \texttt{ifft}(\hat{\boldsymbol{\mathcal{Z}}}, [], 3)$;
return $\boldsymbol{\mathcal{Z}}$

Our next result shows that projection based on slice sampling leads to near optimal estimation in tensor regression, when the covariate tensor has small coherence.

THEOREM 4.7. *Given a target tensor $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and a rank-r tensor of $\boldsymbol{\mathcal{L}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, choose $c \geq \xi r \mu_0(\boldsymbol{\mathcal{V}}) \log(4n_2/\delta)/\epsilon^2$, where ξ is a fixed positive constant and $\delta \in (0, 1]$. Let $\boldsymbol{\mathcal{A}}_c \in \mathbb{R}^{n_1 \times c \times n_2}$ be a tensor of c lateral slices of $\boldsymbol{\mathcal{A}}$ sampled uniformly without replacement, and let $\boldsymbol{\mathcal{L}}_c \in \mathbb{R}^{n_1 \times c \times n_3}$ consist of the corresponding slices of $\boldsymbol{\mathcal{L}}$. Then, with probability at least $1 - \delta - 0.2$,*

$$\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{A}}_c * \boldsymbol{\mathcal{L}}_c^{+} * \boldsymbol{\mathcal{L}}\|_F \leq (1 + \epsilon)\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{L}}^{+} * \boldsymbol{\mathcal{L}}\|_F.$$

Theorem 4.7 implies an estimate guarantee for our t-CX low rank approximation. In the "noiseless" setting, the following consequence of Theorem 4.7 shows the quality of the lateral slice-based low rank approximation.

COROLLARY 4.8. *Given a target tensor $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and a rank-r tensor of $\boldsymbol{\mathcal{L}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, choose $c \geq \xi r \mu_0(\boldsymbol{\mathcal{V}}) \log(n_2) \log(1/\delta)/\epsilon^2$, where ξ is a fixed positive constant. Then, with probability at least $1 - \delta$,*

$$\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{C}}^{+} * \boldsymbol{\mathcal{A}}\|_F \leq (1 + \epsilon)\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{L}}\|_F.$$

4.3. Slice-based low rank approximation: tensor CUR decomposition. Now, we present for the first time the notion of tensor CUR (t-CUR) decomposition based on *circulant algebra*:

DEFINITION 4.9. *(**t-CUR**). Let $\boldsymbol{\mathcal{X}}$ be an $n_1 \times n_2 \times n_3$ tensor. For any given $\boldsymbol{\mathcal{C}}$, an $n_1 \times c \times n_3$ tensor whose slices consist of c lateral slices of the tensor $\boldsymbol{\mathcal{X}}$, and $\boldsymbol{\mathcal{R}}$, an $l \times n_2 \times n_3$ tensor whose slices consist of l horizontal slices of the tensor $\boldsymbol{\mathcal{X}}$, the $n_1 \times n_2 \times n_3$ tensor*

$$\boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{U}} * \boldsymbol{\mathcal{R}},$$

is a lateral-horizontal-based tensor approximation to $\boldsymbol{\mathcal{X}}$ for any $c \times l \times n_3$ tensor $\boldsymbol{\mathcal{U}}$.

Several things should be noted about this definition.

- The t-CUR decomposition is most appropriate as a data analysis tool, when the data consist of one- and/or two modes that are qualitatively different than the remaining ones. In this case, the t-CUR decomposition approximately expresses the original data tensor in terms of a basis consisting of underlying subtensors that are actual data slices and not artificial bases.
- The t-CUR approximation is a t-CX approximation, but one with a very special structure; i.e., every lateral slice of $\boldsymbol{\mathcal{X}}$ can be expressed in terms of the basis provided by $\boldsymbol{\mathcal{C}}$ using only the information contained in a small number of horizontal slices of $\boldsymbol{\mathcal{X}}$ and a low-dimensional encoding tensor.
- In terms of its singular value structure, $\boldsymbol{\mathcal{U}}$ contains the "inverse-of-$\boldsymbol{\mathcal{X}}$" information. For the t-CUR decomposition described in our paper, $\boldsymbol{\mathcal{U}}$ will be a generalized inverse of the intersection between $\boldsymbol{\mathcal{C}}$ and $\boldsymbol{\mathcal{R}}$. More precisely, $\boldsymbol{\mathcal{U}} := \boldsymbol{\mathcal{C}}^{+} * \boldsymbol{\mathcal{X}} * \boldsymbol{\mathcal{R}}^{+}$.

Note that the structural simplicity of a t-CUR tensor decomposition can be provided in the Fourier domain, detailed in Algorithm 5: Algorithm 5 takes as input an $n_1 \times n_2 \times n_3$ tensor $\boldsymbol{\mathcal{A}}$, an $n_1 \times c \times n_3$ tensor $\boldsymbol{\mathcal{C}}$ consisting of a small number of lateral slices of $\boldsymbol{\mathcal{A}}$, and an error parameter ϵ. Letting $\hat{\boldsymbol{\mathcal{C}}} = \hat{\boldsymbol{\mathcal{U}}} * \hat{\Sigma} * \hat{\boldsymbol{\mathcal{V}}}^{\top}$, Algorithm 5 uses the sampling probabilities

$$p_i \geq \frac{\beta}{cn_3} \|\hat{\boldsymbol{\mathcal{U}}}_{i::}^{\top}\|, \qquad \forall i \in [n_1], \qquad \beta \in (0,1], \tag{16}$$

to sample a small number of horizontal slices of $\boldsymbol{\mathcal{A}}$. It returns an $l \times n_2 \times n_3$ tensor $\boldsymbol{\mathcal{R}}$ consisting of a small number of horizontal slices of $\boldsymbol{\mathcal{A}}$ and an $l \times c \times n_3$ tensor $\boldsymbol{\mathcal{S}}$ consisting of the corresponding slices of $\boldsymbol{\mathcal{C}}$.

COROLLARY 4.10. *Given a tensor $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and a rank-r approximation $\boldsymbol{\mathcal{L}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, choose $c \geq \xi r \mu_0(\boldsymbol{\mathcal{V}}) \log(n_2) \log(1/\delta)/\epsilon^2$ with ξ a constant and let $\boldsymbol{\mathcal{C}} \in \mathbb{R}^{n_1 \times c \times n_3}$ be a tensor of c slices of $\boldsymbol{\mathcal{A}}$ sampled uniformly without replacement. Further choose $l \geq \xi c \mu_0(\boldsymbol{\mathcal{U}}_c) \log(n_1) \log(1/\delta')/\epsilon^2$, and let $\boldsymbol{\mathcal{R}} \in \mathbb{R}^{l \times n_1 \times n_2}$ be a tensor of l horizontal slices of $\boldsymbol{\mathcal{A}}$ sampled independently and uniformly without replacement. Then, with probability at least $(1-\delta)(1-\delta'-0.2)$,*

$$\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{U}} * \boldsymbol{\mathcal{R}}\|_F \leq (1+\epsilon)^2 \|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{L}}\|_F$$

5. Tensor completion and robust factorization. Next, we provide a CUR tensor nuclear norm minimization (CUR t-NN) method which solves noisy tensor factorization using a small number of lateral and/or horizontal slices of the underlying tensor and exhibits favorable computational complexity and comes with performance guarantees.

5.1. Related work. *Low rank plus sparse matrix decomposition.* In many image processing and computer vision applications, the given data matrix X can be decomposed as a sum of a low-rank and a sparse component. To that end, Candés et al. proposed Robust PCA [4] to model data matrices generated according to the following mechanism:

$$\min_{L} \operatorname{rank}(L) + \|E\|_0 \qquad s.t. \qquad X = L + E, \tag{17}$$

However, to find the solution of (17) is an NP-hard problem. It is established in [5] that if the rank of L exhibits a certain degree of low-rankness, while E is sparse

Algorithm 5 (t-CUR), a fast Monte-Carlo algorithm for tensor low rank approximation.

Input: $\boldsymbol{\mathcal{A}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, $\boldsymbol{\mathcal{C}} \in \mathbb{R}^{n_1 \times c \times n_3}$ consisting of c lateral slices of $\boldsymbol{\mathcal{A}}$, $p_i \geq 0$, a rank parameter r, and positive integer $l \leq n_1$.
$\hat{\boldsymbol{\mathcal{A}}} \leftarrow \texttt{fft}(\boldsymbol{\mathcal{A}}, [], 3)$;
$\hat{\boldsymbol{\mathcal{C}}} \leftarrow \texttt{fft}(\boldsymbol{\mathcal{C}}, [], 3)$;
Compute sampling probabilities p_i for all $i \in [n_1]$ given by (15);
Initialize $\hat{\boldsymbol{\mathcal{S}}}$ and $\hat{\boldsymbol{\mathcal{D}}}$ to the all zeros tensors.
$t = 1$;
for $i = 1, \ldots, n_1$ **do**
  Pick i with probability $\min\{1, lp_i\}$;
  if i is picked **then**
    $\hat{\boldsymbol{\mathcal{S}}}_{it1} = 1, \quad \hat{\boldsymbol{\mathcal{S}}}_{it2:n_3} = 0$;
    $\hat{\boldsymbol{\mathcal{D}}}_{tt1} = 1/\min\{1, \sqrt{lp_i}\}, \quad \hat{\boldsymbol{\mathcal{D}}}_{tt2:n_3} = 0$;
    $t = t + 1$
  end if
end for
for $k = 1, \ldots, n_3$ **do**
  $\hat{\boldsymbol{\mathcal{R}}}_{::k} = \hat{\boldsymbol{\mathcal{D}}}_{::k} \hat{\boldsymbol{\mathcal{S}}}_{::k}^{\top} \hat{\boldsymbol{\mathcal{A}}}_{::k}$;
  $\hat{\boldsymbol{\mathcal{U}}}_{::k} = \left(\hat{\boldsymbol{\mathcal{D}}}_{::k} \hat{\boldsymbol{\mathcal{S}}}_{::k}^{\top} \hat{\boldsymbol{\mathcal{C}}}_{::k}\right)^{+}$;
  $\hat{\boldsymbol{\mathcal{Z}}}_{::k} = \hat{\boldsymbol{\mathcal{C}}}_{::k} \hat{\boldsymbol{\mathcal{U}}}_{::k} \hat{\boldsymbol{\mathcal{R}}}_{::k}$;
end for
$\boldsymbol{\mathcal{Z}} \leftarrow \texttt{ifft}(\hat{\boldsymbol{\mathcal{Z}}}, [], 3)$;
return $\boldsymbol{\mathcal{Z}}$

enough, then the formulation of (17) can be relaxed into a convex problem of the form:

$$\min_{L} \|L\|_* + \|E\|_1 \qquad s.t. \qquad X = L + E, \tag{18}$$

This model implicitly assumes that the underlying data structure lies in a single low-rank subspace. However, in many applications (e.g. image classifications) it is more likely that the data are obtained from a union of multiple subspaces, and hence recovery of the structure based on the above decomposition would be inaccurate. In order to segment the data into their respective subspaces, one needs to compute an affinity matrix that encodes the pairwise affinities between data vectors. Liu [23] proposed a more general rank minimization problem, where the data matrix itself is used as the dictionary, resulting in the following convex optimization problem:

$$\min_{L} \|L\|_* + \|E\|_{2,1} \qquad s.t. \qquad X = XL + E, \qquad \text{diag}(L) = 0 \tag{19}$$

When the subspaces are *globally* independent, the data are noiseless and sampling is sufficient, Liu et al. [23] show that the optimal solution, denoted by L^*, to the problem given by 19 corresponds to the widely used Shape Iteration Matrix (SIM) [8]. The latter is a "block-diagonal" affinity matrix that indicates the true segmentation of the data. To handle data corrupted by noise, the popular Low Rank Representation (LRR) introduced in [23] adopts a regularized formulation that introduces an extra penalty term to fit the noise component. Further, after obtaining the self-representation matrix L, the affinity matrix C is usually constructed as

$C = \frac{1}{2}(|L| + |L^\top|)$, where $|\cdot|$ represents the absolute operator. Then, the obtained affinity matrix C will be processed through a spectral clustering algorithm [28] to produce the final clustering result.

It is established in [36, 33] that combining sparse and low-rank regularization can improve the performance of image classification. The basic objective function of this combination [36] is as follows:

$$\min_{L} \lambda_1\|L\|_* + \lambda_2\|L\|_1 + \lambda_3\|E\|_\ell \qquad s.t. \qquad X = DL + E, \qquad \mathrm{diag}(Z) = 0 \tag{20}$$

where $\lambda_1, \lambda_2, \lambda_3$ are tuning parameters and $\|E\|_\ell$ indicates different norms suitable for different types of noise components; for example, the squared Frobenius norm for Gaussian noise and the ℓ_1 norm for random spiked noise (corruption of the data). Equation (20) is similar to the objective functions in [33, 12], where a detailed explanation of the formulation given in (20) is also provided.

5.1.1. Low rank tensor decomposition. Based on the new computable tensor SVD, i.e, t-SVD, the tensor nuclear norm [30] is used to replace the tubal rank for low-rank tensor recovery (from incomplete/corrupted tensors) by solving the following convex program,

$$\min_{\boldsymbol{\mathcal{L}}} \|\boldsymbol{\mathcal{L}}\|_{\circledast} \qquad s.t. \qquad \|P_\Omega(\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{L}})\| \le \Delta, \tag{21}$$

where $\|\boldsymbol{\mathcal{L}}\|_{\circledast}$ denotes the tensor nuclear norm (see Section 2 for the definition) and

$$(P_\Omega(\boldsymbol{\mathcal{X}})_{ij} = \boldsymbol{\mathcal{X}}_{ij}, \quad \text{if} \quad (i,j) \in \Omega \quad \text{and} \quad (P_\Omega(\boldsymbol{\mathcal{X}}))_{ij} = 0 \quad \text{otherwise.}$$

Lu et al. [25] extended Robust PCA [4] to the third order tensor based on t-SVD and proposed the following convex problem:

$$\min_{\boldsymbol{\mathcal{L}},\boldsymbol{\mathcal{E}}} \|\boldsymbol{\mathcal{L}}\|_{\circledast} + \lambda\|\boldsymbol{\mathcal{E}}\|_1 \qquad s.t. \qquad \|P_\Omega(\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{L}} - \boldsymbol{\mathcal{E}})\|_F \le \Delta, \tag{22}$$

This model implicitly assumes that the underlying data structure is a single low-rank subspace. When the data is drawn from a union of multiple subspaces, which is common in image classification, the recovery may be inaccurate. Xie et. al[34] extended LRR based subspace clustering to multi-view by employing the rank sum of different mode unfloding to constrain the subspace coefficient tensor, resulting in the following convex optimization problem:

$$\min_{\boldsymbol{\mathcal{L}}} \|\boldsymbol{\mathcal{L}}\|_{\circledast} + \lambda\|\boldsymbol{\mathcal{E}}\|_{2,1} \qquad s.t. \qquad \|P_\Omega(\boldsymbol{\mathcal{X}} - \boldsymbol{\mathcal{X}} * \boldsymbol{\mathcal{L}} - \boldsymbol{\mathcal{E}})\|_F \le \Delta, \tag{23}$$

Equation (23) is similar to the objective functions in [17, 29] and the detail explanation of (23) can also be found there.

5.2. Proposed algorithm. Next, we propose an effective algorithm for large scale noisy tensor decomposition. Our proposal, called *CUR Tensor Nuclear Norm* (CUR t-NN) is an algorithm that extends Algorithm 5 to the tensor completion problems.

Algorithm 6 outlines the main steps. Note that, depending on the application, computing the sampling probabilities p^c and p^l may or may not be necessary. For the convenience of error analysis, we present an implementation of Algorithm 6 in the Fourier domain. This allows us to apply results from the tensor projection independently to each frontal slice.

Algorithm 6 `CUR t-NN`, CUR tensor nuclear norm minimization.

1: $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, positive integer c and l.
2: Sample c lateral slice of $\mathcal{X}$ based on probabilities $p_j^c, j \in [n_2]$, satisfying

$$p_j^c = \frac{\|\mathcal{X}_{:j:}\|_F}{\|\mathcal{X}\|_F},$$

and construct tensor $\mathcal{C} \in \mathbb{R}^{n_1 \times c \times n_3}$.
3: Sample l horizontal slice of $\mathcal{X}$ based on probabilities $p_i^l, i \in [n_1]$, satisfying

$$p_i^l = \frac{\|\mathcal{C}_{i::}\|_F}{\|\mathcal{C}\|_F},$$

and construct tensor $\mathcal{R} \in \mathbb{R}^{l \times c \times n_3}$.
4: Solve tensor completion problem using $P_\Omega(\mathcal{C})$, and $P_\Omega(\mathcal{R})$ and determine the low rank approximations $\tilde{\mathcal{C}}$ and $\tilde{\mathcal{R}}$;
5: Let $\tilde{\mathcal{U}} = \tilde{\mathcal{W}}^+$, where $\tilde{\mathcal{W}}$ is the $l \times c \times n_3$ tensor formed by sampling the corresponding l horizontal slices of $\tilde{\mathcal{C}}$.
6: $\tilde{\mathcal{L}} = \tilde{\mathcal{C}} * \tilde{\mathcal{U}} * \tilde{\mathcal{R}}$;
7: **return** $\tilde{\mathcal{L}}$

5.3. Running Time of CUR t-NN. Algorithm 6 significantly reduces the per-iteration complexity of nuclear norm minimization problems. Indeed, in each iteration, a base tensor nuclear norm minimization algorithm requires $O(n_1 n_2 n_3 log(n_3))$ flops to transform to the Fourier domain, $O(n_1 n_2 n_3 \min(n_1, n_2))$ flops for the t-SVD computation and factorization, and $O(n_1 n_2 n_3 \min(n_1, n_2))$ flops to back to the original domain. On the other hand, Algorithm 6 only requires the corresponding $O(lcn_3 log(n_3))$, $O(lcn_3 \min(l, c))$ and $O(lcn_3 log(n_3))$ flops. Further, Algorithm 6 can be implemented without storing the data tensor $\mathcal{X}$ and can be advantageous when $r \ll \min(n_1, n_2)$ which is common in real data sets. Since the *t-product* and *t-project* of 3-way tensors (see Algorithmw 1 and 3) reduce to the standard matrix product when the third dimension is 1, the `CUR t-NN` reduces to a modified version of matrix nuclear norm minimization algorithm [26] whenever $n_3 = 1$.

5.4. Theoretical guarantees. Now, we show that Algorithm 6 exhibits high probability recovery guarantees comparable to those of the base algorithm. Our first result bounds the μ_0 and μ_1-coherence of a randomized tensor in terms of the coherence of the full tensor.

LEMMA 5.1. *Let $\mathcal{L}_c$ be a tensor of c slices of a rank-r tensor $\mathcal{L} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ sampled uniformly without replacement. If $c \geq \xi r \mu_0(\mathcal{V}) \log(n_2) \log(1/\delta)/\epsilon^2$, where ξ is a positive constant, then with probability at least $1 - \delta/n_2$,*

$$\begin{aligned}
\text{rank}_t(\mathcal{L}_c) &= \text{rank}_t(\mathcal{L}); \\
\mu_0(\mathcal{U}_{\mathcal{L}_c}) &= \mu_0(\mathcal{U}); \\
\mu_0(\mathcal{V}_{\mathcal{L}_c}) &\leq \frac{1}{1-\epsilon/2}\mu_0(\mathcal{V}); \\
\mu_1(\mathcal{L}_c) &\leq \frac{r}{1-\epsilon/2}\mu_0(\mathcal{U})\mu_0(\mathcal{V}),
\end{aligned}$$

Our next theorem provides a bound on the estimation error of `CUR t-NN` used in combination with an arbitrary base algorithm.

THEOREM 5.2. *Under the notation of Algorithms 6, let*

$$c \geq \xi r \mu_0(\boldsymbol{\mathcal{V}}) \log(n_2) \log(2/\delta)/\epsilon^2, \qquad l \geq \xi c \mu_0(\boldsymbol{\mathcal{U}}_{\mathfrak{C}}) \log(n_1) \log(4/\delta')/\epsilon^2,$$

where ξ is a constant and $\delta, \delta' \in (0, 1]$. Let $\boldsymbol{\mathcal{C}}_0$ and $\boldsymbol{\mathcal{R}}_0$ be the corresponding lateral and horizontal sub-tensors of exact solution $\boldsymbol{\mathcal{L}}_0$. If $\boldsymbol{\mathcal{L}}_0$ is (μ, r)-coherent then with probability at least $(1-\delta)(1-\delta'-0.2)$

- *$\boldsymbol{\mathcal{C}}_0$ and $\boldsymbol{\mathcal{R}}_0$ are $(\frac{r\mu^2}{1-\epsilon/2}, r)$-coherent, and*
- $\|\boldsymbol{\mathcal{L}}_0 - \tilde{\boldsymbol{\mathcal{L}}}\|_F \leq (2+3\epsilon)\sqrt{\|\boldsymbol{\mathcal{C}}_0 - \tilde{\boldsymbol{\mathcal{C}}}\|_F^2 + \|\boldsymbol{\mathcal{R}}_0 - \tilde{\boldsymbol{\mathcal{R}}}\|_F^2}$.

6. Performance Evaluation. In this section, we investigate the efficiency of the proposed randomized algorithms, especially the `CUR t-NN`, on the real data sets.

We have implemented the `CUR t-NN` algorithm along with the following algorithms in the MATLAB R2015b environment on a PC with a 1.8 GHz processor and 6 GB RAM memory and double precision format:

- **EXACT NN**, the matrix completion on each frame of the tensor [5];
- **RPCA NN**, the robust matrix completion on each frame of the tensor [4];
- **E-TUCKER NN**, the TUCKER based tensor completion [24];
- **R-TUCKER NN**, the robust TUCKER based tensor completion [14];
- **EXACT t-NN**, the t-SVD based tensor completion [37];
- **RPCA t-NN**, the robust t-SVD based tensor completion [25];

All the algorithms are being terminated either the relative square error (RSE),

$$\text{RSE} := \frac{\|\boldsymbol{\mathcal{L}}_0 - \tilde{\boldsymbol{\mathcal{L}}}\|_F}{\|\boldsymbol{\mathcal{L}}_0\|_F} \leq 10^{-3},$$

or the number of iterations and CPU times exceed 1,000 and 20 minutes, respectively.

6.1. Video completion. We first compare the `CUR t-NN` to the listed competing methods for video data completion from randomly missing entries. We randomly sampled 50% entries from the video. The result is shown in Table 1. It can be seen that the `CUR t-NN` outperforms its competitors in terms of CPU running time and accuracy.

Completion Approach	**RSE**	**Time(s)**
`EXACT NN` [5]	0.1001	687
`E-TUCKER NN`[24]	0.0900	718
`EXACT t-NN` [37]	0.0715	695
CUR t-NN	**0.0740**	**115**

Table 1: RSE of tensor completion results for the basketball video.

6.2. Cropped Yale Face B Dataset. We apply the `CUR t-NN` on the Cropped Yale B dataset to evaluate the accuracy of the proposed low-rank representations, as well as the computation time of the `CUR t-NN`. The proposed algorithms are demonstrated on an application to facial recognition. The datasets for the experiments are a subset of the Cropped Extended Yale Face Dataset B. The Cropped Yale B dataset

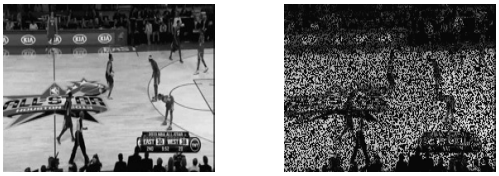

(a)

(b)

Fig. 6: The 20th frame of tensor completion result on a basketball video. **(a) Left**: The original video. **(a)Right**: Sampled video (50% sampling rate). **(b) Left**: `EXACT t-NN` [37]. **(b) Right**: `CUR t-NN`.

contains 1140 images that has the first 30 possible illuminations of 38 different people. Each image has 192×168 pixels in a grayscale range. This is collected into a $192 \times 1140 \times 168$ tensor, and this tensor is denoted by B.

Completion Approach	**RSE**	**Time(s)**
`RPCA NN` [4]	0.0056	417
`R-TUCKER NN`[14]	0.0034	513
`RPCA t-NN` [25]	0.0021	495
CUR t-NN	**0.0026**	**136**

Table 2: RSE of tensor robust completion results for Cropped Yale B dataset

Table 2 shows the accuracy of the `CUR t-NN` in comparison with other algorithms. As shown in Figure 7 and Table 2, `CUR t-NN` estimated nearly the same face model as the `RPCA t-NN` in a small fraction of time. Indeed, the `CUR t-NN` exhibited the RSE 0.0026, quite small given pixels. The associated running time was reduced from 495s for `RPCA t-NN` to 136s for `CUR t-NN`.

7. Conclusion. In this paper, we discussed the disadvantages and limitations of the deterministic tensor principle component analysis. The CUR t-NN, we design and analysed combines the advantages of randomization for dimensionality reduction and applied it to large scale noisy tensor decompositions. We also extended the randomized CX and CUR method to the third order tensors based on slice selection and provide a basic version algorithm (t-CX), as well as a more general version algorithm (t-CUR with subspace iterations). The differences between the CUR t-NN and the deterministic approaches highlight some of the main advantages of this work: i) CUR t-NN can be used in combination with any underlying tensor factorization algorithm,

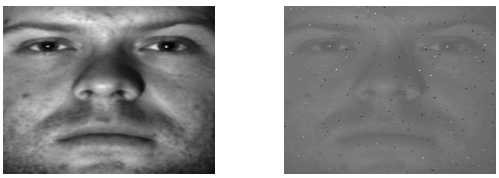

(a)

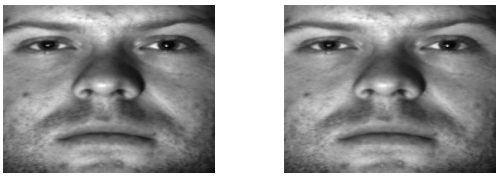

(b)

Fig. 7: The 10th frame of tensor completion result on a Yale Face B dataset. **(a) Left**: The original frame. **(a)Right**: Noisy image (20% pixels corrupted). **(b) Left**: `SVD t-NN` [25]. **(b) Right**: `CUR t-NN`.

ii) CUR t-NN is trivially has a linear running time , and iii) CUR t-NN provably maintains the recovery guarantees of its base algorithm, even in the presence of noise.

8. Acknowledgments.

REFERENCES

[1] A. Anandkumar, R. Ge, D. J. Hsu, S. M. Kakade, and M. Telgarsky, *Tensor decompositions for learning latent variable models.*, Journal of Machine Learning Research, 15 (2014), pp. 2773–2832.

[2] T. Bouwmans, A. Sobral, S. Javed, S. K. Jung, and E.-H. Zahzah, *Decomposition into low-rank plus additive matrices for background/foreground separation: A review for a comparative evaluation with a large-scale dataset*, arXiv preprint arXiv:1511.01245, (2015).

[3] K. Braman, *Third-order tensors as linear operators on a space of matrices*, Linear Algebra and its Applications, 433 (2010), pp. 1241–1253.

[4] E. J. Candès, X. Li, Y. Ma, and J. Wright, *Robust principal component analysis?*, Journal of the ACM (JACM), 58 (2011), p. 11.

[5] E. J. Candès and B. Recht, *Exact matrix completion via convex optimization*, Foundations of Computational mathematics, 9 (2009), pp. 717–772.

[6] J. D. Carroll and J.-J. Chang, *Analysis of individual differences in multidimensional scaling via an n-way generalization of eckart-young decomposition*, Psychometrika, 35 (1970), pp. 283–319.

[7] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. A. Phan, *Tensor decompositions for signal processing applications: From two-way to multiway component analysis*, IEEE Signal Processing Magazine, 32 (2015), pp. 145–163.

[8] J. P. Costeira and T. Kanade, *A multibody factorization method for independently moving objects*, International Journal of Computer Vision, 29 (1998), pp. 159–179.

[9] A. Deshpande, L. Rademacher, S. Vempala, and G. Wang, *Matrix approximation and projective clustering via volume sampling*, in Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, Society for Industrial and Applied Mathematics, 2006, pp. 1117–1126.

[10] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast monte carlo algorithms for matrices i: Approximating matrix multiplication*, SIAM Journal on Computing, 36 (2006), pp. 132–157.

[11] P. Drineas, M. W. Mahoney, and S. Muthukrishnan, *Relative-error cur matrix decompositions*, SIAM Journal on Matrix Analysis and Applications, 30 (2008), pp. 844–881.
[12] P. Foggia, G. Percannella, and M. Vento, *Graph matching and learning in pattern recognition in the last 10 years*, International Journal of Pattern Recognition and Artificial Intelligence, 28 (2014), p. 1450001.
[13] D. Hsu, S. M. Kakade, T. Zhang, et al., *Tail inequalities for sums of random matrices that depend on the intrinsic dimension.*
[14] B. Huang, C. Mu, D. Goldfarb, and J. Wright, *Provable low-rank tensor recovery*, (2014).
[15] I. Jolliffe, *Principal component analysis*, Wiley Online Library, 2002.
[16] E. Karahan, P. A. Rojas-Lopez, M. L. Bringas-Vega, P. A. Valdes-Hernandez, and P. A. Valdes-Sosa, *Tensor analysis and fusion of multimodal brain images*, Proceedings of the IEEE, 103 (2015), pp. 1531–1559.
[17] E. Kernfeld, S. Aeron, and M. Kilmer, *Clustering multi-way data: a novel algebraic approach*, arXiv preprint arXiv:1412.7056, (2014).
[18] M. E. Kilmer, K. Braman, N. Hao, and R. C. Hoover, *Third-order tensors as operators on matrices: A theoretical and computational framework with applications in imaging*, SIAM Journal on Matrix Analysis and Applications, 34 (2013), pp. 148–172.
[19] M. E. Kilmer and C. D. Martin, *Factorization strategies for third-order tensors*, Linear Algebra and its Applications, 435 (2011), pp. 641–658.
[20] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM review, 51 (2009), pp. 455–500.
[21] P. M. Kroonenberg, *Three-mode principal component analysis: Theory and applications*, vol. 2, DSWO press, 1983.
[22] F. G. Kuruvilla, P. J. Park, and S. L. Schreiber, *Vector algebra in the analysis of genome-wide expression data*, Genome biology, 3 (2002), pp. research0011–1.
[23] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, *Robust recovery of subspace structures by low-rank representation*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 171–184.
[24] J. Liu, P. Musialski, P. Wonka, and J. Ye, *Tensor completion for estimating missing values in visual data*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 208–220.
[25] C. Lu, J. Feng, Y. Chen, W. Liu, Z. Lin, and S. Yan, *Tensor robust principal component analysis: Exact recovery of corrupted low-rank tensors via convex optimization.*
[26] L. W. Mackey, M. I. Jordan, and A. Talwalkar, *Divide-and-conquer matrix factorization*, in Advances in Neural Information Processing Systems, 2011, pp. 1134–1142.
[27] M. Mohri and A. Talwalkar, *Can matrix coherence be efficiently and accurately estimated?*
[28] A. Y. Ng et al., *On spectral clustering: Analysis and an algorithm*, (2002).
[29] X. Piao, Y. Hu, J. Gao, Y. Sun, and Z. Lin, *A submodule clustering method for multi-way data by sparse and low-rank representation*, arXiv preprint arXiv:1601.00149, (2016).
[30] O. Semerci, N. Hao, M. E. Kilmer, and E. L. Miller, *Tensor-based formulation and nuclear norm regularization for multienergy computed tomography*, IEEE Transactions on Image Processing, 23 (2014), pp. 1678–1693.
[31] A. Smilde, R. Bro, and P. Geladi, *Multi-way analysis: applications in the chemical sciences*, John Wiley & Sons, 2005.
[32] M. A. O. Vasilescu and D. Terzopoulos, *Multilinear analysis of image ensembles: Tensorfaces*, in European Conference on Computer Vision, Springer, 2002, pp. 447–460.
[33] Y.-X. Wang, H. Xu, and C. Leng, *Provable subspace clustering: When lrr meets ssc*, in Advances in Neural Information Processing Systems, 2013, pp. 64–72.
[34] Y. Xie, D. Tao, W. Zhang, and L. Zhang, *Multi-view subspace clustering via relaxed l_1-norm of tensor multi-rank*, arXiv preprint arXiv:1610.07126, (2016).
[35] J. Zhang, A. K. Saibaba, M. Kilmer, and S. Aeron, *Divide-and-conquer matrix factorization*, in arXiv preprint arXiv:1609.07086, 2016, pp. 1–26.
[36] Y. Zhang, Z. Jiang, and L. S. Davis, *Learning structured low-rank representations for image classification*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 676–683.
[37] Z. Zhang and S. Aeron, *Exact tensor completion using t-svd*, arXiv preprint arXiv:1502.04689, (2015).
[38] Z. Zhang, G. Ely, S. Aeron, N. Hao, and M. Kilmer, *Novel methods for multilinear data completion and de-noising based on tensor-svd*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 3842–3849.

SUPPLEMENTARY MATERIALS

.1. Proof of Theorem 3.1.

Proof. We prove that for any set of probabilities $\{p_j\}_{j=1}^{n_2}$ the following holds:

$$\mathbf{E}\|\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}\|_F^2 \leq \frac{1}{c}\sum_{j\in\Lambda}\frac{\|\boldsymbol{\mathcal{A}}_{:j:}\|_F^2\|\boldsymbol{\mathcal{B}}_{j::}\|_F^2}{p_j}. \tag{24}$$ ∎

Let I_j, $j \in [n_2]$ be the indicator variable that is set to 1 if the j-th lateral slices of $\boldsymbol{\mathcal{A}}$ and the j-th horizontal slices of $\boldsymbol{\mathcal{B}}$ are sampled and is set to 0 otherwise. Using the Algorithm 2, if $I_j = 1$, we scale both the j-th lateral slice of $\boldsymbol{\mathcal{A}}$ and the j-th horizontal slice of $\boldsymbol{\mathcal{B}}$ by $1/\sqrt{\min\{1, cp_j\}}$. Thus,

$$\begin{aligned}
\|\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}\|_F^2 &= \|(F_{n_3}\otimes I)\text{circ}(\boldsymbol{\mathcal{A}})(F_{n_3}^{-1}\otimes I)(F_{n_3}\otimes I)\text{unfold}(\boldsymbol{\mathcal{B}}) \\
&\quad - (F_{n_3}\otimes I)\text{circ}(\boldsymbol{\mathcal{C}})(F_{n_3}^{-1}\otimes I)(F_{n_3}\otimes I)\text{unfold}(\boldsymbol{\mathcal{R}})\|_F^2 \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\|\hat{\boldsymbol{\mathcal{A}}}_{::k}\hat{\boldsymbol{\mathcal{B}}}_{::k} - \hat{\boldsymbol{\mathcal{C}}}_{::k}\hat{\boldsymbol{\mathcal{R}}}_{::k}\|_F^2 \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\|\hat{\boldsymbol{\mathcal{A}}}_{::k}\hat{\boldsymbol{\mathcal{B}}}_{::k} - \hat{\boldsymbol{\mathcal{A}}}_{::k}\hat{\boldsymbol{\mathcal{S}}}_{::k}\hat{\boldsymbol{\mathcal{D}}}_{::k}\hat{\boldsymbol{\mathcal{D}}}_{::k}\hat{\boldsymbol{\mathcal{S}}}_{::k}^{\top}\hat{\boldsymbol{\mathcal{B}}}_{::k}\|_F^2 \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\|\sum_{j=1}^{n_2}\left(1 - \frac{I_j}{\min\{1, cp_j\}}\right)\hat{\boldsymbol{\mathcal{A}}}_{:jk}\hat{\boldsymbol{\mathcal{B}}}_{j:k}\|_F^2,
\end{aligned}$$

where the first equality follows from the unitary invariance of Frobenius norm.

If $\min\{1, cp_j\} = 1$, then $I_j = 1$ with probability 1, and $1 - I_j/\min\{1, cp_j\} = 0$. Therefore, we take the expectation of the right sides of (25) on the set of indices $\Lambda = \{j \in [n_2] : cp_j < 1\} \subseteq [n_2]$,

$$\begin{aligned}
\mathbf{E}\left[\|\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}\|_F^2\right] &= \frac{1}{n_3}\sum_{k=1}^{n_3}\mathbf{E}\left[\|\sum_{j\in\Lambda}\left(1 - \frac{I_j}{cp_j}\right)\hat{\boldsymbol{\mathcal{A}}}_{:jk}\hat{\boldsymbol{\mathcal{B}}}_{j:k}\|_F^2\right] \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\mathbf{E}\left[\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\left(\sum_{j\in\Lambda}\left(1 - \frac{I_j}{cp_j}\right)\hat{\boldsymbol{\mathcal{A}}}_{:jk}\hat{\boldsymbol{\mathcal{B}}}_{j:k}\right)^2_{i_1 i_2}\right] \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\mathbf{E}\left[\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\left(\sum_{j\in\Lambda}\left(1 - \frac{I_j}{cp_j}\right)\hat{\boldsymbol{\mathcal{A}}}_{i_1 jk}\hat{\boldsymbol{\mathcal{B}}}_{ji_2 k}\right)^2\right].
\end{aligned}$$

By multiplying out the right hand side, it follows that

$$\begin{aligned}
\mathbf{E}\left[\|\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}\|_F^2\right] &= \frac{1}{n_3}\sum_{k=1}^{n_3}\mathbf{E}\left[\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\sum_{j_1\in\Lambda}\sum_{j_2\in\Lambda}\hat{p}_{i_1 i_2 j_1 j_2}\right] \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\sum_{j_1\in\Lambda}\sum_{j_2\in\Lambda}\mathbf{E}\left[\hat{p}_{i_1 i_2 j_1 j_2}\right],
\end{aligned}$$

where $\hat{p}_{i_1 i_2 j_1 j_2} = \left(1 - \frac{I_{j_1}}{cp_{j_1}}\right)\left(1 - \frac{I_{j_2}}{cp_{j_2}}\right)\hat{\boldsymbol{\mathcal{A}}}_{i_1 j_1 k}\hat{\boldsymbol{\mathcal{B}}}_{j_1 i_2 k}\hat{\boldsymbol{\mathcal{A}}}_{i_1 j_2 k}\hat{\boldsymbol{\mathcal{B}}}_{j_2 i_2 k}$.

Now, using the fact that $j \in [\Lambda]$, $\mathbf{E}[1 - I_j/cp_j] = 0$ and $\mathbf{E}[(1 - I_j/cp_j)^2] = (1/cp_j) - 1 \leq 1/cp_j$, we obtain

$$
\begin{aligned}
\mathbf{E}\left[\|\boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{B}} - \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{R}}\|_F^2\right] &= \frac{1}{n_3}\sum_{k=1}^{n_3}\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\sum_{j\in\Lambda}\mathbf{E}(1 - I_j/cp_j)^2\hat{\boldsymbol{\mathcal{A}}}_{i_1jk}^2\hat{\boldsymbol{\mathcal{B}}}_{ji_2k}^2 \\
&\leq \frac{1}{n_3}\sum_{k=1}^{n_3}\sum_{j\in\Lambda}\frac{1}{cp_j}\sum_{i_1=1}^{n_1}\sum_{i_2=1}^{n_2}\hat{\boldsymbol{\mathcal{A}}}_{i_1jk}^2\hat{\boldsymbol{\mathcal{B}}}_{ji_2k}^2 \\
&= \frac{1}{n_3}\sum_{k=1}^{n_3}\frac{1}{c}\sum_{j\in\Lambda}\frac{\|\hat{\boldsymbol{\mathcal{A}}}_{:jk}\|_2^2\|\hat{\boldsymbol{\mathcal{B}}}_{j:k}\|_2^2}{p_j} \\
&= \frac{1}{c}\sum_{j\in\Lambda}\frac{\|\boldsymbol{\mathcal{A}}_{:j1}\|_2^2\|\boldsymbol{\mathcal{B}}_{j:1}\|_2^2}{p_j} + \cdots + \frac{1}{c}\sum_{j\in\Lambda}\frac{\|\boldsymbol{\mathcal{A}}_{:jn_3}\|_2^2\|\boldsymbol{\mathcal{B}}_{j:n_3}\|_2^2}{p_j} \qquad (25) \\
&= \frac{1}{c}\sum_{j\in\Lambda}\frac{\|\boldsymbol{\mathcal{A}}_{:j:}\|_F^2\|\boldsymbol{\mathcal{B}}_{j::}\|_F^2}{p_j}.
\end{aligned}
$$

Now, equation(11) follows from (24) by using Jensen's inequality and the fact that the sampling probabilities (9) and (10) are defined in the original domain.

.2. Proof of Corollary 3.2.

Proof. Once Theorem 3.1 is substituted for Theorem 7 of [11], the proof of this corollary follows similarly from [13]. ∎

.3. Proof of Lemma 4.5.

Proof. Using Corollary 3.2, we get the following with probability at least $1 - \delta$,

$$
\|\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{V}} - \boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{V}}\|_{\mathrm{F}}^2 \leq \epsilon/2\|\boldsymbol{\mathcal{V}}^\top\|_{\mathrm{F}}^2, \qquad (26)
$$

This further gives us a bound on the singular values of $\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}$, for all $1 \leq i \leq r$ and $1 \leq k \leq n_3$,

$$
\begin{aligned}
|1 - \max_{k\in[n_3]}\sigma_{i,k}^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}})| &= |\max_{k\in[n_3]}\sigma_{i,k}^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{V}}) - \max_{k\in[n_3]}\sigma_{i,k}^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{V}})| \\
&\leq |\sigma^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{V}}) - \sigma^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{V}})| \\
&= \|\boldsymbol{\mathcal{V}} * \boldsymbol{\mathcal{V}}^\top - \boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{V}}\|^2 \\
&\leq \|\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{V}} - \boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{D}} * \boldsymbol{\mathcal{S}}^\top * \boldsymbol{\mathcal{V}}\|_{\mathrm{F}}^2 \\
&\leq \epsilon/2, \qquad (27)
\end{aligned}
$$

where σ is the largest singular value of tensor. Note that the second inequality follows since $\|\cdot\| \leq \|\cdot\|_F$ and the last inequality follows from (26).

Thus, it follows that for all $1 \leq i \leq r$ and $1 \leq k \leq n_3$,

$$
1 - \epsilon/2 \leq \max_{k\in[n_3]}\sigma_{i,k}^2(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}) \leq 1 + \epsilon/2 \qquad (28)
$$

which implies that each tubal singular value of $\boldsymbol{\mathcal{V}}$ is positive, and so $\mathrm{rank}_t(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}}) = \mathrm{rank}_t(\boldsymbol{\mathcal{V}}) = \mathrm{rank}_t(\boldsymbol{\mathcal{L}})$.

To prove the second claim, we use the t-SVD of $\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}$ and note that

$$\begin{aligned}
\|\Omega\| &= \|\left(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}\right)^+ - \left(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}\right)^\top\| \\
&= \|\left(U_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \boldsymbol{\mathcal{V}}^\top_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\right)^+ - \left(\boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \boldsymbol{\mathcal{V}}^\top_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\right)^\top\| \\
&= \|\boldsymbol{\mathcal{V}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \left(\Sigma^{-1}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} - \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\right) \boldsymbol{\mathcal{U}}^\top_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\| \\
&= \| \left(\Sigma^{-1}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} - \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\right) \|. \qquad (29)
\end{aligned}$$

The claim follows since $\boldsymbol{\mathcal{V}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}$ and $\boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}$ are orthogonal tensors.

To prove the third claim, note that

$$\begin{aligned}
(\boldsymbol{\mathcal{L}} * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}})^+ &= \left(\boldsymbol{\mathcal{U}} * \Sigma * \boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}\right)^+ \\
&= \left(\boldsymbol{\mathcal{U}} * \Sigma * \boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \boldsymbol{\mathcal{V}}^\top_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\right)^+ \\
&= \boldsymbol{\mathcal{V}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * (\Sigma * \boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}})^+ * \boldsymbol{\mathcal{U}}^\top. \qquad (30)
\end{aligned}$$

To remove the pseudoinverse in the above derivations, we use the first part. In this case,

$$\begin{aligned}
(\Sigma * \boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}})^+ &= (\Sigma * \boldsymbol{\mathcal{U}}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}})^{-1} \\
&= \Sigma^{-1}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} * \boldsymbol{\mathcal{U}}^\top_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} \Sigma^{-1}. \qquad (31)
\end{aligned}$$

By combining (30) and (31), we get the result.

To prove the last claim, we have that for all $1 \le i \le r$ and $1 \le k \le n_3$,

$$\begin{aligned}
\|\Omega\| &= \|\Sigma^{-1}_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}} - \Sigma_{\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}}\| && \text{from (29)} \\
&= \max_{i,k} \left| \sigma_{i,k}(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}) - \frac{1}{\sigma_{i,k}(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}})} \right| && \text{by definition} \\
&= \max_{i,k} \frac{|\sigma^2_{i,k}(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}}) - 1|}{|\sigma_{i,k}(\boldsymbol{\mathcal{V}}^\top * \boldsymbol{\mathcal{S}} * \boldsymbol{\mathcal{D}})|} && \text{by simple manipulation} \\
&\le \frac{\epsilon/2}{\sqrt{1-\epsilon/2}} && \text{from (27)} \\
&\le \epsilon/\sqrt{2} && \text{if } \epsilon < 1 \Rightarrow \sqrt{1-\epsilon/2} > 1/\sqrt{2}. \square
\end{aligned}$$

.4. Proof of Theorem 4.7.

Proof. Using Proposition 4.6, and the fact that $n_2 \ge r\frac{\mu_0(\boldsymbol{\mathcal{V}})}{n_3}$ we have

$$c \ge \xi r \frac{\mu_0(\boldsymbol{\mathcal{V}})}{n_3} \log(4r\mu_0(\boldsymbol{\mathcal{V}})/n_3\delta)/\epsilon^2 \ge \xi r \log(4r/(\beta\delta))/(\beta\epsilon^2)$$

whenever $\beta \ge n_3/\mu_0(\boldsymbol{\mathcal{V}})$. From (15) with $\beta = n_3/\mu_0(\boldsymbol{\mathcal{V}}) \in (0,1]$ and $p_i = 1/n_2$, we obtain

$$\frac{\beta}{rn_3}\|\hat{\boldsymbol{\mathcal{V}}}^\top_{:i:}\|_F^2 = \frac{\beta}{r}\|\boldsymbol{\mathcal{V}}^\top_{:i:}\|_F^2 \le \frac{\beta}{r}\frac{r}{n_2 n_3}\mu_0(\boldsymbol{\mathcal{V}}) = \frac{1}{n_2} = p_i$$

for all $i \in [n_2]$, by the definition of $\mu_0(\boldsymbol{\mathcal{V}})$.

Since $\boldsymbol{\mathcal{D}} = \boldsymbol{\mathcal{I}}/\sqrt{n_2/c}$, with probability at least $1 - \delta - 0.2$, we have that

$$\begin{aligned}
\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{A}}_c * \boldsymbol{\mathcal{L}}_c^+ * \boldsymbol{\mathcal{L}}\|_F &= \|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{C}} * \boldsymbol{\mathcal{D}} * (\boldsymbol{\mathcal{L}}_c * \boldsymbol{\mathcal{D}})^+ * \boldsymbol{\mathcal{L}}\|_F \\
&\le (1+\epsilon)\|\boldsymbol{\mathcal{A}} - \boldsymbol{\mathcal{A}} * \boldsymbol{\mathcal{L}}^+ * \boldsymbol{\mathcal{L}}\|_F.
\end{aligned}$$

.5. Proof of Corollary 4.8.

Proof. Since $\mathcal{C}^+ * \mathcal{A}$ minimizes $\|\mathcal{A} - \mathcal{C} * \mathcal{X}\|_F$ over all tensor $\mathcal{X} \in \mathbb{R}^{n_1 \times c \times n_3}$, it follows that

$$\|\mathcal{A} - \mathcal{C} * \mathcal{C}^+ * \mathcal{A}\|_F \leq \|\mathcal{A} - \mathcal{C} * \mathcal{L}_{\mathcal{C}}^+ * \mathcal{L}\|_F,$$

Thus, using Theorem 4.7 and Proposition 4.6 we have

$$\|\mathcal{A} - \mathcal{C} * \mathcal{C}^+ \mathcal{A}\|_F \leq (1+\epsilon)\|\mathcal{A} - \mathcal{L}\|_F,$$

The reminder of the proof follows from Theorem 6 in [26]. ∎

.6. Proof of Corollary 4.10.

Proof. Using Theorem 4.7, Proposition 4.6 and Corollary 4.8 it follows that

$$\|\mathcal{A} - \mathcal{C} * \mathcal{U} * \mathcal{R}\|_F \leq (1+\epsilon)\|\mathcal{A} - \mathcal{C} * \mathcal{C}^+ * \mathcal{A}\|_F \leq (1+\epsilon)^2\|\mathcal{A} - \mathcal{L}\|_F$$

with probability at least $(1-\delta)(1-\delta'-0.2)$ by independence. ∎

.7. Proof of Lemma 5.1.

Proof. The first claim follows from Lemma 4.5 with $\beta = n_3/\mu_0(\mathcal{V}) \in (0,1]$ and $p_i = 1/n_2$. Using Lemma 4.5, the second claim follows along similar lines to the proof of Lemma 1 of [27].

To prove the third claim, using Lemma 4.5, assume that $\mathcal{S}^\top * \mathcal{V}$ consists of the first c horizontal slices of $\mathcal{V}$. Then if $\mathcal{L}_c = \mathcal{U} * \Sigma * \mathcal{V}^\top * \mathcal{S}$ has $\text{rank}_t(\mathcal{L}_c) = \text{rank}_t(\mathcal{L}) = r$, the tensor $\mathcal{V}^\top * \mathcal{S}$ must have full tubal rank. Thus we can write

$$\begin{aligned}
\mathcal{L}_c^+ * \mathcal{L}_c &= (\mathcal{U} * \Sigma * \mathcal{V}^\top * \mathcal{S})^+ * \mathcal{U} * \Sigma * \mathcal{V}^\top * \mathcal{S} \\
&= (\Sigma * \mathcal{V}^\top * \mathcal{S})^+ * \mathcal{U}^\top * \mathcal{U} * \Sigma * \mathcal{V}^\top * \mathcal{S} \\
&= (\Sigma * \mathcal{V}^\top * \mathcal{S})^+ * \Sigma * \mathcal{V}^\top * \mathcal{S} \\
&= (\mathcal{V}^\top * \mathcal{S})^+ * \Sigma^\top * \Sigma * \mathcal{V}^\top * \mathcal{S} \\
&= (\mathcal{V}^\top * \mathcal{S})^+ * \mathcal{V}^\top * \mathcal{S} \\
&= \mathcal{V}^\top * \mathcal{S} * (\mathcal{V}^\top * \mathcal{S} * \mathcal{S}^\top * \mathcal{V})^{-1} \mathcal{V}^\top * \mathcal{S},
\end{aligned}$$

where the second and third equalities follow from $\mathcal{U}^\top * \mathcal{U} = \mathcal{I}_c$. The fourth and fifth result from Σ having full tubal rank and $\mathcal{V}$ having full lateral slice rank, and the sixth follows from $\mathcal{S}^\top * \mathcal{V}$ having full horizontal slice rank. Now, denote the right singular vectors of $\mathcal{L}_c$ by $\mathcal{V}_c \in \mathbb{R}^{c \times r \times n_3}$. Define $\mathring{\mathbf{e}}_{i,c}$ as the ith lateral slice of $\mathcal{I}_c$ and $\mathring{\mathbf{e}}_{i,n_2}$ as the ith lateral slice of $\mathcal{I}$. Then we have,

$$\begin{aligned}
\mu_0(\mathcal{V}_c) &= \frac{cn_3}{r} \max_{1 \leq i \leq c} \|\mathcal{V}_c^\top * \mathring{\mathbf{e}}_{i,c}\|_F^2 \\
&= \frac{cn_3}{r} \max_{1 \leq i \leq c} \text{trace}\{\mathring{\mathbf{e}}_{i,c}^\top * \mathcal{L}_c^+ * \mathcal{L}_c * \mathring{\mathbf{e}}_{i,c}\} \\
&= \frac{cn_3}{r} \max_{1 \leq i \leq c} \text{trace}\{\mathring{\mathbf{e}}_{i,c}^\top * (\mathcal{V}^\top * \mathcal{S})^+ * \mathcal{V}^\top * \mathcal{S} * \mathring{\mathbf{e}}_{i,c}\} \\
&= \frac{cn_3}{r} \max_{1 \leq i \leq c} \text{trace}\{\mathring{\mathbf{e}}_{i,c}^\top * \mathcal{S}^\top * \mathcal{V} * \mathcal{W}^{-1} * \mathcal{V}^\top * \mathcal{S} * \mathring{\mathbf{e}}_{i,c}\} \\
&= \frac{cn_3}{r} \max_{1 \leq i \leq c} \text{trace}\{\mathring{\mathbf{e}}_{i,n_2}^\top * \mathcal{V} * \mathcal{W}^{-1} * \mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2}\},
\end{aligned}$$

where $\mathcal{W} = \mathcal{V}^\top * \mathcal{S} * \mathcal{S}^\top * \mathcal{V}$ and the final equality follows from $\mathcal{V}^\top * \mathcal{S} * \mathring{\mathbf{e}}_{i,c} = \mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2}$ for all $1 \leq i \leq c$.

Now, we have

$$\begin{aligned}\mu_0(\mathcal{V}_{\mathcal{L}_e}) &= \frac{cn_3}{r}\max_{1\le i\le c}\text{trace}\{\mathring{\mathbf{e}}_{i,n_2}^\top * \mathcal{V} * \mathcal{W}^{-1} * \mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2}\}\\ &= \frac{cn_3}{r}\max_{1\le i\le c}\text{trace}\{\mathcal{W}^{-1} * \mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2} * \mathring{\mathbf{e}}_{i,n_2}^\top * \mathcal{V}\}\\ &\le \frac{cn_3}{r}\|\mathcal{W}^{-1}\|^2\max_{1\le i\le c}\|\mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2} * \mathring{\mathbf{e}}_{i,n_2}^\top * \mathcal{V}\|_{\circledast}^2,\end{aligned}$$

where the last inequality follows form Hölder's inequality.

Since $\mathcal{V}^\top * \mathring{\mathbf{e}}_{i,n_2} * \mathring{\mathbf{e}}_{i,n_2}^\top * \mathcal{V}$ has tubal rank one, using the definition of μ_0-coherence, we have

$$\mu_0(\mathcal{V}_{\mathcal{L}_e}) \le \frac{c}{n_2}\|\mathcal{W}^{-1}\|^2\mu_0(\mathcal{V}).$$

Now, using (28), we have that $\|\mathcal{W}^{-1}\|^2 \le \frac{n_2}{c(1-\epsilon/2)}$. Thus, it follows that

$$\mu_0(\mathcal{V}_{\mathcal{L}_e}) \le \mu_0(\mathcal{V})/(1-\epsilon/2).$$

To prove the last claim under Lemma 4.5, we note that

$$\begin{aligned}\mu_1(\mathcal{L}_c) &= \frac{n_1cn_3^2}{r}\max_{\substack{1\le i\le n_1\\ 1\le j\le c}}\|\mathring{\mathbf{e}}_{i,n_1}^\top * \mathcal{U}_c * \mathcal{V}_c^\top * \mathring{\mathbf{e}}_{j,c}\|_F^2\\ &\le \frac{n_1cn_3^2}{r}\max_{1\le i\le n_1}\|\mathcal{U}_c^\top * \mathring{\mathbf{e}}_{i,n_1}\|_F^2\max_{1\le j\le c}\|\mathcal{V}_c^\top * \mathring{\mathbf{e}}_{j,c}\|_F^2\\ &\le \frac{r}{(1-\epsilon/2)}\mu_0(\mathcal{U})\mu_0(\mathcal{V}).\end{aligned}$$

.8. Proof of Theorem 5.2.

Proof. To prove Theorem 5.2, let $\mathcal{L}$ denotes the solution obtained by a base Algorithm. Under the notion of algorithm 6, we partition $\mathcal{L}$ as:

$$\mathcal{L} = \begin{bmatrix}\mathcal{C}_1 & \mathcal{R}_2\\ \mathcal{C}_2 & \mathcal{L}_{0,22}\end{bmatrix}, \quad \text{where} \quad \mathcal{C} = \begin{bmatrix}\mathcal{C}_1\\ \mathcal{C}_2\end{bmatrix}, \quad \mathcal{R} = \begin{bmatrix}\mathcal{R}_1 & \mathcal{R}_2\end{bmatrix}$$

and $\mathcal{L}_{0,22} \in \mathbb{R}^{(n_1-l)\times(n_2-c)\times n_3}$ is the bottom right subtensor of $\mathcal{L}_0$. Define $W(\mathcal{L})$ as the event

$$\|\mathcal{L}-\tilde{\mathcal{L}}\|_F \le (1+\epsilon)^2\|\mathcal{L}_0-\mathcal{L}\|_F. \tag{32}$$

∎

Using Corollary 4.10, $W(\mathcal{L})$ holds with probability at least $(1-\delta/2)(1-\delta/4-0.2)$. Moreover, $A(\mathcal{C})$ and $A(\mathcal{R})$ hold with probability at least $1-\delta/(2n_1)$ and $1-\delta/(4n_2)$ respectively by Lemma 5.1. Thus, for $n_2 \ge 2$ and $\delta \le 0.8$,

$$\begin{aligned}\mathbf{Pr}(W(\mathcal{L})\bigcap A(\mathcal{C})\bigcap A(\mathcal{R})) &\ge 1-\mathbf{Pr}(W(\mathcal{L})^c)-\mathbf{Pr}(A(\mathcal{C})^c)-\mathbf{Pr}(A(\mathcal{R})^c)\\ &\ge 1-(1-(1-\delta/2)(1-\delta/4-0.2))-\delta/(2n_1)-\delta/(4n_2)\\ &\ge (1-\delta)(1-\delta-0.2).\end{aligned}$$

Now, using (32), we have

$$\|\mathcal{L}_0-\tilde{\mathcal{L}}\|_F \le \|\mathcal{L}_0-\mathcal{L}\|_F+\|\mathcal{L}-\tilde{\mathcal{L}}\|_F \le (2+2\epsilon+\epsilon^2)\|\mathcal{L}_0-\mathcal{L}\|_F \le (2+3\epsilon)\|\mathcal{L}_0-\mathcal{L}\|_F,$$

when $W(\mathcal{L})$ holds, by the triangle inequality.